arXiv:1610.08580v1 [stat.ME] 27 Oct 2016

# A Generalized Approach to Power Analysis for Estimation of Local Average Treatment Effects

Kirk Bansak*

kbansak@stanford.edu

October 2016

Abstract

This paper introduces a new approach to power analysis in the context of estimating a local average treatment effect (LATE), where the study subjects exhibit noncompliance with treatment assignment. As a result of distributional complications in the LATE context, compared to the simple ATE context, there is currently no standard method of power analysis for the LATE estimator. Moreover, existing methods and commonly used substitutes—which include intent-to-treat (ITT) and scaled ATE power analyses—require specifying generally unknown variance parameters and/or rely upon strong and unrealistic assumptions, thus providing unreliable guidance on the power of the LATE estimator. This paper develops a framework that uses standardized effect sizes to place bounds on the LATE power formula whereby distributional parameters need not be specified nor assumed. Instead, in addition to the effect size, sample size, and error tolerance parameters, the only other parameter that must be specified by the researcher is the first-stage effect. Additional conditions can also be introduced to further narrow the bounds on the power calculation. The result is a generalized approach to power analysis in the LATE context that is simple to implement. Software is provided to implement the methods.

Keywords: experimental design, local average treatment effects, noncompliance, principal stratification, statistical power

*Department of Political Science, Stanford University, 616 Serra Street, Stanford, CA 94305-6044.

I thank Avi Acharya, Jens Hainmueller, Dominik Hangartner, Andy Hall, and Teppei Yamamoto for helpful comments.

# I. Introduction

Power analysis has long been recognized as a vital study design tool (Cohen, 1962). Running simple power analyses provides researchers with concrete and reliable information to inform their budgetary requirements, choose a sample size, and form reasonable expectations on the magnitude of treatment effects they will be able to detect. This helps researchers avoid an eventuality in which a study has failed to produce meaningful findings not because there is nothing interesting to find but rather due to insufficient power to overcome fundamentally noisy data. The results of a power analysis can also help researchers avoid running certain studies altogether if the costs are simply too prohibitive in light of the probability of successful detection of a meaningful effect. Yet various fields of research are replete with studies that have failed to report power analyses and implemented drastically under-powered designs (Tversky and Kahneman, 1971; Tsang et al., 2009). Indeed, many researchers' (and funders') time, energy, and money has been put at serious risk by neglect of power analyses in the early stages of research design.

In practice, however, power analyses can often be challenging to properly implement. In the case of the simple research design involving a randomized treatment and the goal of estimating an average treatment effect (ATE), a complication is the need, *a priori*, for estimates of (or informal guesses about) the variance of the outcome variable (Bloom, 2006; Duflo et al., 2007). While previous studies and/or existing data can often help to inform these variance estimates, there often do not exist any data that is recent or closely related enough to serve as a useful benchmark, particularly when a researcher is interested in a novel outcome variable or a new population of interest. This complication in a simple ATE power analysis is well-known among applied researchers. Fortunately, there also exist fixes for this problem, as will be described later.

Less well-known, however, is the extent to which similar complications become magnified in the context of estimating a local average treatment effect (LATE), where the study subjects exhibit noncompliance with treatment assignment. Due to the existence of multiple principal strata, distinct distributional behavior across those strata, and the focus on local identification of the average treatment effect, the number of distributional parameters that impact power vastly proliferates in the LATE context. In fact, the power of LATE hypothesis tests not only depends upon the rate of compliance, with lower compliance resulting in lower power, but is also impacted by the different means and variances of all of the principal strata as well as the proportions of subjects belonging to each principal stratum.

As a result, there is currently no standard method of power analysis in the LATE context. This state of affairs is particularly concerning given how common noncompliance is in many research environments, including field experiments, clinical trials, and randomized controlled trials (RCTs) using encouragement designs. These types of studies also tend to be among the most expensive, generating strong incentives for well-calibrated power analyses. Yet existing methods require specifying distributional assumptions that are difficult to make and/or come with hidden, implicit assumptions about the various principal strata that are unlikely to reflect the reality of one's applied data. In addition, recognizing the complexity of LATE power analysis, some researchers settle for performing scaled ATE or intent-to-treat (ITT) power analyses even when their ultimate estimand of interest is the LATE—a precarious practice given that, as shown later, the results of scaled ATE and ITT power analyses can diverge dramatically from the results of a LATE power analysis.

This paper introduces a new approach to LATE power analysis. Specifically, by using a particular type of standardized (rather than absolute) LATE effect size, this paper shows how bounds can be placed on the power formula whereby neither variance components nor patterns of noncompliance and heterogeneity need to be specified. Instead, in addition to the effect size, sample size, and error tolerance parameters, the only other parameter that must be specified by the researcher is the first-stage effect. In contrast, nine other parameters that affect test power need not be specified. As usual, the effect size and sample size parameters can be isolated, presenting formulas for minimum detectable effect sizes and necessary sample sizes. Additional assumptions can also be made to further narrow the bounds on the power calculation. The result is a generalized approach to power analysis in the LATE context that is simultaneously conservative, disciplined, and simple to implement. The approach can also be used as a power test for fuzzy regression discontinuity designs (Hahn et al., 2001) that use the instrumental variable (IV) estimator in the discontinuity window around the threshold, as well as for quasi-experiments that apply the IV framework to observational data.

## II. Power Analysis for Average Treatment Effects

Let $Y_0$ and $Y_1$ denote the potential outcomes of the dependent variable under control ($D = 0$) and treatment ($D = 1$) conditions, respectively. Further, let $\hat{\bar{Y}}_0$ and $\hat{\bar{Y}}_1$ denote the averages of the observed outcome $Y$ for those units assigned to the control and treatment conditions, respectively, which constitute unbiased estimates of the mean values of $Y_0$ and $Y_1$ given random assignment of the treatment. As a result, $\hat{\delta} = \hat{\bar{Y}}_1 - \hat{\bar{Y}}_0$ is an unbiased estimate of the sample

average treatment effect (ATE), the true value of which is $\delta$.

As shown elsewhere (Bloom, 2006; Duflo et al., 2007), given asymptotic normality of the ATE estimator, power analysis then proceeds with the following equation:

$$\Phi\left(-c^* + \frac{\delta}{\sqrt{V_N}}\right) + \Phi\left(-c^* - \frac{\delta}{\sqrt{V_N}}\right) = 1 - \beta$$

where $\Phi(\bullet)$ denotes the standard normal cumulative distribution function, $\delta$ denotes a hypothesized true ATE value, $V_N$ denotes the sampling variance of the ATE estimator ($\hat{\delta}$), $1 - \beta$ denotes the power to correctly reject the hypothesis test ($\beta$ denotes the type-II error rate), and $c^*$ denotes the critical value corresponding to the tolerable type-I error rate ($\alpha$) and hypothesis test type. For the standard two-tailed test of the null hypothesis of zero difference in means, $c^* = \Phi^{-1}(1 - \frac{\alpha}{2})$. Conventionally, $V_N \equiv \frac{Var(Y_1)}{N_1} + \frac{Var(Y_0)}{N_0}$, which is conservative given complete randomization of $N$ units with a pre-determined number $N_1$ assigned to treatment and $N_0 = N - N_1$ assigned to control (Imbens and Rubin, 2015, Chapter 6). The standard power formula displayed above can be rearranged depending upon the specific parameter for which the analyst wants to solve. Common parameters to solve for include $\delta$, so that an analyst can recover a minimum detectable effect, and $N$, so that the analyst can recover the necessary sample size.

### A. ATE Power Analysis Complications

In order to use the power formula above, the analyst must specify $V_N$, which requires explicitly or implicitly specifying the variances of $Y_0$ and $Y_1$. This requirement presents a possibly serious practical complication. While previous studies and/or existing data can often help to inform these variance specifications, there often do not exist any data that is recent or closely related enough to serve as a useful benchmark, particularly in the case where a researcher is interested in a novel outcome variable or new population of interest. In the case of a binary outcome variable, this problem is partially alleviated by the ability to easily choose a conservative, upper-bound variance. In the case of continuous outcome variables, however, such a mathematical convenience does not exist. One might devise an idea, based on theoretical expectations, about what a conservative variance might look like. However, in the very plausible case that these expectations are inaccurate, too high a guess will lead to an overpowered study while too low a guess will lead to an unsuccessful study. In both cases, the researcher's resources are at risk of being wasted.

Furthermore, lack of reliable knowledge about the potential outcome variances also complicates a researcher's ability to specify meaningful treatment effect values, which is already a

subjective task even with known variances. While it is useful to think about effect magnitudes in absolute terms, often times it is difficult to predict what a "meaningfully large" absolute effect magnitude would be, particularly with less tangible types of outcome variables. Furthermore, from a policy or project evaluation standpoint, it may not be useful to conceptualize the utility or cost-efficiency of an intervention in absolute terms—i.e. without reference to the distribution of the outcome variable of interest. As a result, nonexistent or ill-informed variance estimates doubly hinder proper power analysis by preventing the researcher from specifying reasonable effect magnitudes.

### B. Standardized Effect Sizes

"Effect sizes" have long been established as the standard solution to this problem in the ATE context (Cohen, 1988; Bloom, 2006; Duflo et al., 2007). In cases where the variances are unknown and/or absolute effect magnitudes are difficult to interpret, a common recommendation is to employ the effect size $\frac{\delta}{\sigma}$—rather than the absolute effect $\delta$—where $\sigma$ is the standard deviation of a reference outcome distribution. In other words, effect sizes are measures of treatment effects that are standardized with reference to the distribution of the outcome variable. By employing effect sizes, the result is that the variance terms in the power formula drop out, thus obviating the inconvenient need to estimate or guess variance values. In addition, there exist general benchmarks for what constitutes a small, medium, and large effect size (Cohen, 1988), and meta-analyses within individual fields of study have enabled researchers to develop discipline-specific guidance on effect size significance (Lipsey, 1990, Chapter 3).

## III. Power Analysis in the LATE Context

### A. Local Average Treatment Effect (LATE) Setup

In many studies, the subjects exhibit noncompliance: some units assigned to the treatment condition do not take the treatment, and/or some units assigned to the control condition do take the treatment. This problem is pervasive across many research settings—including field experiments, clinical trials, and RCTs using encouragement designs—as subjects often cannot be forced to take the treatment, and some subjects find ways to access the treatment even when not assigned to it (Gerber and Green, 2012). In the presence of noncompliance, the ATE generally cannot be identified. However, it is possible to identify a local average treatment effect (LATE), which is the ATE for a particular subset of the population: the compliers, or those who take the treatment when assigned to the treatment and do not otherwise (Angrist

et al., 1996). The compliers are one of three subsets, or principal strata, in the population. The second is the never-takers, or those who never take the treatment regardless of their assignment. The third is the always-takers, or those who always take the treatment regardless of their assignment.[1] In the case of one-sided noncompliance, the LATE is the ATE for the treated (given no always-takers) or the ATE for the untreated (given no never-takers).

The identification and estimation framework for the LATE is as follows. Let $Y$ denote the observed outcome, $D$ the treatment uptake, and $Z$ the randomized treatment assignment/encouragement. Assume $D$ and $Z$ are both binary. Let the following equations denote the relationships between the three variables:

$$D - E[D] = \pi \cdot (Z - E[Z]) + \nu_1 \tag{1}$$

$$Y - E[Y] = \tau \cdot (D - E[D]) + \nu_2 \tag{2}$$

$$Y - E[Y] = \gamma \cdot (Z - E[Z]) + \nu_3 \tag{3}$$

Note that in this notation, $\tau$, $\pi$, and $\gamma$ correspond to average effects, rather than constant parameters. While $\tau$ denotes the treatment effect, it is not identified by equation (2) due to the endogeneity of $D$ and effect heterogeneity. Instead, given a set of standard assumptions[2] laid out by Angrist et al. (1996), $\tau$ is identified as the LATE by $\frac{\gamma}{\pi}$ as in the classic instrumental-variables (IV) approach, where $\pi$ is the first-stage effect of treatment assignment on treatment uptake and $\gamma$ is the intent-to-treat (ITT) effect. As shown by Imbens and Angrist (1994), the LATE can be estimated consistently as follows:

$$\widehat{LATE} = \frac{\widehat{Cov}(Y, Z)}{\widehat{Cov}(D, Z)}$$

The asymptotic variance of the estimator is:

$$V_N^{\widehat{LATE}} = \frac{E[\nu_2^2\{Z - E[Z]\}^2]}{NCov^2(D, Z)}$$

[1] A fourth stratum, defiers, are those who take the treatment when not assigned to it but do not take the treatment when assigned to it; however, identification of the LATE requires assuming the non-existence of defiers.

[2] The assumptions include (1) the stable unit treatment value assumption (SUTVA), (2) the treatment assignment $Z$ is random, (3) the exclusion restriction, whereby $Z$ only has an impact on $Y$ via $D$, (4) a nonzero average causal effect of the treatment assignment $Z$ on treatment uptake $D$, and (5) monotonicity, which given binary $Z$ and $D$ means the non-existence of defiers, or units that take the treatment when not assigned to it and do not take the treatment when assigned to it.

## B. Literature Review

In contrast to the ATE context, there is no standard method of power analysis for the LATE estimator.

Bloom (2006) and Duflo et al. (2007) have laid out standard methods for the power analysis of ATEs, and furthermore, both discuss issues presented by study subject noncompliance. However, neither of these studies develops a formal power analysis for the LATE context. Bloom discusses noncompliance issues, but does not address power within the context of noncompliance. Duflo et al. go further in discussing and partially formalizing the relationship between power and noncompliance. In fact, they even offer a formula for computing minimum detectable effects in the presence of noncompliance based on a simple scaling of the standard ATE formula, the result of which follows from the ATE estimator variance being divided by the compliance rate squared. The problem with this solution, which Duflo et al. do not make explicit but has been shown elsewhere (Baiocchi et al., 2014), is that a number of strong and unrealistic assumptions are required for this scaling of the standard ATE power to be (approximately) equal to the LATE power. Specifically, it must be the case that the never-takers have the same mean outcome value as the untreated compliers, the always-takers have the same mean outcome value as the treated compliers, and all groups have the same outcome variance. Furthermore, the formula also requires specifying the outcome variance, a parameter that may not be available to the researcher.

There have also been some studies in the epidemiology literature on Mendelian randomization focused directly on power analysis of IV treatment effects (Pierce et al., 2011; Freeman et al., 2013; Brion et al., 2013). In particular, Freeman et al. (2013) and Brion et al. (2013) both introduce power formulas for IV treatment effects. However, there are two major limitations to the approaches taken by these studies. First, they require specifying a number of variance components, on which a researcher may not have good preexisting knowledge or priors. Second, they proceed from a classic IV perspective and hence neglect the extent to which these variance components depend upon the distinct distributional behavior of the principal strata.

For instance, the formulas presented by Freeman et al. and Brion et al. both require specifying $Var(D)$.[3] This presents a challenge given noncompliance with the treatment assignment, as $Var(D)$ is a function of both the first stage effect $\pi$ (which is also the compliance rate) and

[3] $Var(D)$ enters both formulas directly as well as through the need to specify $\rho_{DZ}$ (the correlation between $D$ and $Z$), which can only be mapped from a hypothetical first-stage effect $\pi$ by specifying both $Var(D)$ and $Var(Z)$.

the proportion of always-takers versus never-takers. In other words, to choose an informative value of $Var(D)$, one must specify not only a hypothetical compliance rate but also the precise pattern of noncompliance. In addition, Freeman et al. also require specifying $Var(Y|D)$, while Brion et al. require specifying the biased asymptotic value of the least squares estimator of the effect of D on Y. Both of these are parameters that a researcher is unlikely to know prior to performing a study, when the power analysis is needed. Finally, both formulas also rely on a simplifying but generally incorrect assumption that $E[\nu_2^2\{Z - E[Z]\}^2] = \sigma_2^2 Var(Z)$ in formulating the IV estimator's variance.

In some research contexts, there may be sufficient preexisting data and previous studies to inform the selection of all the parameters that Freeman et al. and Brion et al. require. Mendelian randomization may indeed be one such area, which explains why Freeman et al. and Brion et al. propose formulas that include these various parameters. However, in many study contexts in the social sciences, medicine, public health, and other fields, the researcher will lack solid estimates or priors on one or more of these parameters. In such contexts, the formulas offered by Freeman et al. and Brion et al. cannot be used reliably.

In sum, very little scholarship has focused on developing power analysis techniques for the LATE context. As a result, this is an area in the applied methodological literature that remains underdeveloped. An earlier study by Jo (2002) recognizes and describes the complications that multiple principal strata introduce to the power analysis context, as will be explained further in the next subsection. However, Jo's study does not offer any concrete methodological solutions. Recognizing the complexity of LATE power analysis, some researchers settle for performing a scaled ATE or ITT power analysis even when their ultimate estimand of interest is the LATE. This is a precarious practice given that the results of scaled ATE and ITT power analyses can diverge dramatically from the results of a LATE power analysis. Yet this is often the best a researcher can do given the lack of any reliable power analysis tools that can deal with a number of complications that arise in the LATE context. The following subsection will describe these complications in greater detail.

### C. LATE Power Analysis Complications

A number of unique characteristics of the LATE context make specifying reasonable variance inputs and meaningful effect sizes a more complex task than in the ATE context.

A first complication that has been recognized for some time relates to the local identification of the LATE estimator (Jo, 2002). The possibility of heterogeneous potential outcome distributions across the three principal strata (compliers, never-takers, and always-takers) combined

with the possibility of different patterns of noncompliance—i.e. the proportions of never-takers and always-takers—creates several additional parameters that must be specified explicitly or implicitly in order to perform a LATE power analysis. In other words, as will be illustrated later, the unique characteristics of the LATE context lead to the proliferation of parameters that affect the power.

This first problem leads to a second complication in terms of being able to specify standardized effect sizes in such a way that the variance components of the power analysis drop out. As in the ATE context, the possibility of employing effect sizes suggests a direction for eliminating distributional assumptions. However, whereas in the ATE context the fix is fairly simple and hence enables the analyst to minimize the number of assumptions that must be made, such a fix has been elusive in the LATE context. The result is that existing power analyses in the LATE context, whether analytic or simulation-based, have inconveniently required (explicit or implicit) distributional assumptions that may not match the reality of the data that will eventually be collected.

Finally, a third complication relates to how to even define effect sizes in the LATE context. The problem is that it is unclear what the correct reference variance for the effect sizes should be. In the ATE context, it is natural to think of the variance of $Y_0$, or the pooled variance of $Y_0$ and $Y_1$, as the reference variance. In the LATE context, however, there is a proliferation of potential outcomes as a result of imperfect compliance and multiple principal strata, and it is unclear which should and should not theoretically be of interest in the reference distribution. For instance, one might think of $Y_{D=0}$, $Y_{D=1}$, $Y_{Z=0}$, or $Y_{Z=1}$ as possible elements of the reference variance, where $D$ denotes treatment uptake and $Z$ denotes treatment assignment/encouragement. However, in contrast to the ATE context, neither $Y_{D=0}$ nor $Y_{Z=0}$ constitute a natural state of the world. The distribution of $Y_{D=0}$ is no longer the natural state of no intervention because there is self-selection into and out of the treatment by some units. Furthermore, $Y_{Z=0}$ also does not represent any meaningful state in the absence of intervention for the same reasons. Finally, while a seemingly natural reference distribution in the LATE case could theoretically be represented by $Y_{D=0}$ for the compliers, there is no way to observe this reference distribution because of the inability to distinguish untreated compliers from never-takers. Similarly, $Y_{D=1}$ for the compliers cannot be observed because of the inability to distinguish treated compliers from always-takers.

### D. Parameters in LATE Power Analysis

In general, we may separate power analysis parameters into three groups: error tolerance parameters, investigation parameters, and distribution parameters. The error tolerance parameters are $\alpha$ (type-I error tolerance) and $\beta$ (type-II error tolerance), where $\beta$ is a parameter only in the case where we are solving for a different parameter rather than calculating the power $(1 - \beta)$. The investigation parameters are sample size and effect magnitude/size. These are the parameters of fundamental economic interest that motivate the use of a pre-study power analysis. Finally, the distribution parameters are the parameters that characterize the distribution(s) of the population(s) of interest. In contrast to the tolerance parameters, which are selected by convention or on a discretionary basis, and the investigation parameters, which the researcher seeks to learn about in order to make research design decisions, the distribution parameters are matters of inconvenience. While they are (usually) not of strict interest to the researcher, the distribution parameters can have a dramatic impact on the calibration of the power analysis, and they must be specified at values that are known or believed to reflect reality in order for the power analysis to be truly informative.

As described above, in the standard ATE context using absolute effect magnitudes, the power formula requires specification of the variance of the ATE estimator, which in turn requires two distribution parameters: the outcome variances of both the treatment and control groups. In addition, by employing standardized effect sizes, these distribution parameters can be dispensed with. In contrast, there are many more distribution parameters in the LATE context. If we consider the first-stage effect $\pi$, which is the compliance rate and hence also the proportion of compliers in the population, to be an investigation parameter in the LATE context, then there are nine distribution parameters that affect the variance of the LATE estimator and hence also affect the power. The reason for this proliferation of parameters has to do with heterogeneity across the principal strata and the local focus of the LATE, which leads the LATE estimator variance to be affected not only by the variances of the three principal strata, but also their relative mean locations and their relative proportions. The resulting distribution parameters are: (1) the complier control condition mean, (2) the complier control condition variance, (3) the complier treatment condition variance, (4) the never-taker (control) condition mean, (5) the never-taker (control) condition variance, (6) the always-taker (treatment) condition mean, (7) the always-taker (treatment) condition variance, (8) the proportion of never-takers in the population, and (9) the proportion of always-takers

in the population.[4]

The reason that all of these parameters have an impact on power is that they collectively determine a number of individual variance and covariance components that factor into the LATE estimator's sampling variance. This property of the LATE estimator is illustrated in the Supplementary Materials (SM) Appendix B, which presents the results of a series of simulations illustrating the power of the LATE estimator as the distributional characteristics of the principal strata are varied. The simulations (Tables B1 and B2) show how, in addition to the investigation parameters, all nine of the distribution parameters affect the power of the LATE estimator.

In addition, the simulations also demonstrate that common practices used as ad hoc simplified approaches to LATE power analysis—namely ITT and scaled ATE power analyses—do not provide reliable substitutes for a formal LATE power analysis. As shown in Appendix B (Table B3), and highlighted in previous work by Jo (2002), the test powers of the LATE and ITT estimators change at different rates as the simulation specifications are altered. Depending upon the behavior of the principal strata, the power of the LATE estimator may be higher or lower than the power of the ITT estimator, demonstrating that a power test for the ITT cannot serve as a replacement for a LATE power analysis.

Appendix B also presents comparisons of the test power of the LATE estimator with the power implied by scaling a standard ATE power analysis. As described by previous studies (Duflo et al., 2007; Baiocchi et al., 2014), one approach to computing the power of the LATE estimator is by performing a simple ATE power analysis but scaling the variance of the standard ATE estimator, specifically dividing it by the compliance rate squared. However, the approximate equality of these two quantities is only achieved given the strong assumptions that (a) the never-takers have the same mean outcome value as the untreated compliers, (b) the always-takers have the same mean outcome value as the treated compliers, and (c) all groups have the same outcome variance. If any of those assumptions are violated, the true power of the LATE estimator can diverge dramatically from the power implied by this scaled ATE power analysis. Appendix B (Table B4) demonstrates this result, illustrating how the scaled ATE power analysis provides extremely unreliable guidance on the power of the LATE

[4]The treatment condition mean for the compliers need not be specified since it is simply the sum of the compliers' control condition mean and the LATE, both of which are specified. Further, treatment (control) condition parameters need not be specified for the never-takers (always-takers) because the treatment (control) condition never manifests in the data for the never-takers (always-takers) by definition. Finally, it should be noted that the sum of (8), (9), and $\pi$ must be one, removing a degree of freedom in the specification of distribution parameters.

estimator.

# IV. A Generalized Approach to Power Analysis for Local Average Treatment Effects

This section introduces a new approach to LATE power analysis. The approach entails two innovations. First, it introduces effect sizes that are measured with reference to the total variance of the observed outcome $Y$. In other words, the effect size will be the effect of the treatment on the outcome (for compliers) in terms of standard deviations of the observed outcome across the entire sample. In the LATE context, as described in subsection III.C, the presence of noncompliance means that the distributions of specific treatment groups (by treatment assignment or treatment uptake) and specific potential outcomes do not represent conceptually clear or natural reference distributions, as they would in the case of full compliance. Furthermore, the one set of potential outcomes that would arguably provide a natural reference—$Y_D$ for the compliers—is not fully observable in the LATE context. In contrast, the simplicity, conceptual clarity, and complete observability of the full distribution of the observed outcome make it a much more natural reference distribution than the other alternatives.

The second innovation and the major contribution of the method introduced in this paper is in showing how, by employing these effect sizes, bounds can be placed on the power formula whereby neither variance components nor patterns of heterogeneity and noncompliance need to be specified. Instead, in addition to the effect size, sample size, and error tolerance parameters, the only other parameter that must be specified by the researcher is the first-stage effect. In other words, only tolerance and investigation parameters must be specified; the analyst need not specify nor even make assumptions about any of the nine distribution parameters described earlier.

## A. Deriving a Modified Power Formula

For all of the results that follow, the standard LATE assumptions will be maintained. Similar to the ATE context, we also invoke asymptotic normality of the LATE estimator. Hence, the power formula begins as:

$$\Phi\left(-c^* + \frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}\right) + \Phi\left(-c^* - \frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}\right) = 1 - \beta$$

where $\tau = \frac{\gamma}{\pi}$ denotes the LATE. We further assume that assignment to the treatment is randomized with equal probability of being assigned to the treatment and control conditions. That is, $P(Z = 1) = 0.5$. Note that this assumption will be relaxed later.

In order to derive a LATE power formula that does not require specifying any distribution parameters, or parameters that depend on them, the objective is to re-express $\frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}$ exclusively in terms of investigation parameters. To begin this process, the variance of the LATE estimator, $V_N^{\widehat{LATE}} = \frac{E[\nu_2^2\{Z-E[Z]\}^2]}{NCov^2(D,Z)}$, can be simplified as follows:

**Lemma 1:** *Given* $P(Z = 1) = 0.5$,[5]

$$V_N^{\widehat{LATE}} = \frac{E[\nu_2^2]}{0.25N\pi^2}$$

It is inconvenient to work with $E[\nu_2^2]$. Given imperfect compliance with the treatment assignment in the LATE context, and hence selection into (or out of) the treatment for some subjects, equation (2) is incompletely specified and $\nu_2$ is not orthogonal to $D$—the reason that equation (2) alone does not allow for identification of a causal effect. As a result, as has been observed elsewhere (Angrist et al., 1996, p. 446), $\nu_2$ is not an intrinsically meaningful disturbance vector. By extension, $E[\nu_2^2]$ is not a substantively meaningful term, nor will it be useful for achieving our objective.

Instead, as will become clear, working in terms of variances conditional upon $Z$ will facilitate a solution to the problem. With this in mind, the variance of the LATE estimator can be re-expressed as follows:

**Lemma 2:** *Given* $P(Z = 1) = 0.5$,

$$V_N^{\widehat{LATE}} = \frac{\sigma_3^2 + \tau^2\sigma_1^2 - 2\tau\sigma_{13}}{0.25N\pi^2}$$

*where* $\sigma_1^2 = Var(\nu_1)$, $\sigma_3^2 = Var(\nu_3)$, *and* $\sigma_{13} = Cov(\nu_1, \nu_3)$.

It will also be inconvenient to work with $\sigma_{13}$, and thus it is useful to eliminate it from the expression. Furthermore, as discussed later, the value and even sign of $\sigma_{13}$ cannot be known without making additional assumptions. However, $\sigma_{13}$ can be eliminated by deriving the following strict bounds:

[5]Proofs of all lemmas and theorems can be found in the SM, Appendix A.

**Lemma 3:** *Assuming without loss of generality that* $\tau > 0$,

$$\frac{\sigma_3^2 + \tau^2\sigma_1^2 - 2\tau\sigma_1\sigma_3}{0.25N\pi^2} \leq \widehat{V_N^{LATE}} \leq \frac{\sigma_3^2 + \tau^2\sigma_1^2 + 2\tau\sigma_1\sigma_3}{0.25N\pi^2}$$

Hence, we have the following bounds for $\widehat{V_N^{LATE}}$, re-expressed to reflect that $\tau = \frac{\gamma}{\pi}$:

$$\widehat{V_N^{LATE}} \textit{ bounds}: \quad \frac{\sigma_3^2 + \frac{\gamma^2}{\pi^2}\sigma_1^2 \pm 2\frac{\gamma}{\pi}\sigma_1\sigma_3}{0.25N\pi^2}$$

As discussed, the effect of interest is defined in standardized terms with reference to the full distribution of the observed outcome variable:

**Definition 1:** *Define the effect size of interest as:*

$$\kappa = \frac{\tau}{\sqrt{Var(Y)}}$$

By focusing on the effect size, $\kappa$ will now take the place of $\tau$ as one of the three investigation parameters, along with $\pi$ and $N$. Accordingly, $\tau$ and $\gamma$ can be re-expressed:

**Lemma 4:** *Given* $\kappa = \frac{\tau}{\sqrt{Var(Y)}}$,

$$\tau = \kappa\sqrt{\sigma_3^2 + 0.25\gamma^2}$$

*and*

$$\gamma = \sqrt{\frac{\kappa^2\sigma_3^2\pi^2}{1 - 0.25\kappa^2\pi^2}}$$

Finally, given Definition 1 and Lemmas 1-4, the following bounds can be put on $\frac{\tau}{\sqrt{\widehat{V_N^{LATE}}}}$:

**Theorem 1:** *Given* $P(Z = 1) = 0.5$*, Definition 1, and Lemmas 1-4:*

$$\frac{\tau}{\sqrt{\widehat{V_N^{LATE}}}} \textit{ bounds}: \quad \sqrt{\frac{\kappa^2(1 + 0.25(\frac{\kappa^2\pi^2}{1-0.25\kappa^2\pi^2}))0.25N\pi^2}{1 + \frac{\kappa^2\sigma_1^2}{1-0.25\kappa^2\pi^2} \pm \frac{2\kappa\sigma_1}{\sqrt{1-0.25\kappa^2\pi^2}}}}$$

Notably, this re-expression achieves the elimination of $\sigma_3^2$, leaving only one remaining term that is not an investigation parameter: $\sigma_1^2$. However, since $D$ is binary, a practical and conservative re-expression of $\sigma_1^2$ can be undertaken by setting $\sigma_1^2$ to its largest possible value as a function of $\pi$.

**Lemma 5:**

$$\sigma_1^2 \leq \left(0.5 - \frac{\pi}{2}\right)\left(0.5 + \frac{\pi}{2}\right)$$

Using the results from Lemma 5 and Theorem 1, the final bounds on $\frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}$ for implementation in the power analysis are as follows:

**Theorem 2:** *Given* $P(Z = 1) = 0.5$*, and setting* $\sigma_1^2$ *to its largest possible value,* $(0.5-\frac{\pi}{2})(0.5+\frac{\pi}{2})$*, the following bounds then contain the true lower bound and a slightly conservative upper bound:*

$$\frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}\ final\ bounds: \quad \sqrt{\frac{\kappa^2(1+0.25(\frac{\kappa^2\pi^2}{1-0.25\kappa^2\pi^2}))0.25N\pi^2}{1+\frac{\kappa^2(0.5-\frac{\pi}{2})(0.5+\frac{\pi}{2})}{1-0.25\kappa^2\pi^2} \pm \frac{2\kappa\sqrt{(0.5-\frac{\pi}{2})(0.5+\frac{\pi}{2})}}{\sqrt{1-0.25\kappa^2\pi^2}}}}$$

As can be seen, these final bounds contain only the three investigation parameters: $\kappa$, $\pi$, and $N$. Furthermore, by providing a strict lower bound, they also provide a strict lower bound for the power. The result is a new set of bounds, free of distributional assumptions, that can serve as the basis for a generalized power analysis for the LATE estimator.

## B. Calculating Power

Recall that the power formula is $\Phi\left(-c^* + \frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}\right) + \Phi\left(-c^* - \frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}\right) = 1 - \beta$. We can now plug in the formula in Theorem 2 for $\frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}$. Values of $\kappa$, $\pi$, and $N$—as well as a type-I error tolerance $\alpha$ to calculate $c^*$—can then be selected in order to calculate the power, $1 - \beta$, of the two-sided test of the hypothesis that the LATE is zero.

As can be seen, there do not exist any variance terms in this modified power formula. As a result, it provides bounds that capture any distributional patterns among all three principal strata. In other words, this modified method of power analysis makes no assumptions, either explicitly or implicitly, about any of the nine distribution parameters described earlier.

In addition, as with standard power analyses, the modified power formula can be rearranged to solve for the investigation parameters. Instead of calculating power based on a specific effect size ($\kappa$), compliance rate ($\pi$), and sample size ($N$), it can be more useful to select a specific power level and solve for one of the other parameters by fixing the rest. Of particular interest in this case should be $\kappa$, solving for which will yield the minimum detectable effect size (MDES),[6] and $N$, solving for which will yield the required sample size.

[6]This term is borrowed from Bloom (2006), who used it in the ATE context as an extension of minimum detectable effects measured in absolute terms (Bloom, 1995).

## C. Solving for the Minimum Detectable Effect Size

Without loss of generality, assume $\tau > 0$. Then, for reasonably high levels of power—which includes conventional power levels, such as 0.8—the second term in the power formula is negligible, simplifying the formula to:

$$\Phi\left(-c^* + \frac{\tau}{\sqrt{\widehat{V_N^{LATE}}}}\right) = 1 - \beta$$

Recalling that $c^* = \Phi^{-1}(1 - \frac{\alpha}{2})$ for a two-sided test, this can then be re-expressed as:

$$\frac{\tau}{\sqrt{\widehat{V_N^{LATE}}}} = \Phi^{-1}\left(1 - \frac{\alpha}{2}\right) + \Phi^{-1}(1 - \beta)$$

Let $M = \Phi^{-1}(1 - \frac{\alpha}{2}) + \Phi^{-1}(1 - \beta)$, which is called the "multiplier." As a result, the power formula leads to:

$$\sqrt{\frac{\kappa^2(1 + 0.25(\frac{\kappa^2\pi^2}{1-0.25\kappa^2\pi^2}))0.25N\pi^2}{1 + \frac{\kappa^2(0.5-\frac{\pi}{2})(0.5+\frac{\pi}{2})}{1-0.25\kappa^2\pi^2} \pm \frac{2\kappa\sqrt{(0.5-\frac{\pi}{2})(0.5+\frac{\pi}{2})}}{\sqrt{1-0.25\kappa^2\pi^2}}}} = M$$

$\kappa$ can then be isolated such that the MDES can be computed as a function of the other parameters. The resulting bounds on the MDES are as follows:

$$\kappa_{Low} = 2\sqrt{\frac{Q^2 + 4Q\pi^2 - 4\sqrt{Q^3\pi^2(1-\pi^2)}}{Q^2 + 16Q\pi^4 - 8Q\pi^2 + 16\pi^4}}$$

$$\kappa_{High} = 2\sqrt{\frac{Q^2 + 4Q\pi^2 + 4\sqrt{Q^3\pi^2(1-\pi^2)}}{Q^2 + 16Q\pi^4 - 8Q\pi^2 + 16\pi^4}}$$

where

$$Q = \frac{4M^2}{N}$$

$\kappa_{High}$ represents the conservative bound.

## D. Solving for the Sample Size

Instead of isolating $\kappa$ as above, $N$ can be isolated in order to solve for the required sample size. Continuing to assume without loss of generality that $\tau > 0$, the bounds on the required sample size are:

$$N_{Low} = \frac{(\kappa^2M^2 - 2\kappa^2M^2\pi^2 + 4M^2)\sqrt{4 - \kappa^2\pi^2} - (8\kappa M^2 - 2\kappa^3M^2\pi^2)\sqrt{(1-\pi)(1+\pi)}}{\kappa^2\pi^2\sqrt{4 - \kappa^2\pi^2}}$$

$$N_{High} = \frac{(\kappa^2 M^2 - 2\kappa^2 M^2 \pi^2 + 4M^2)\sqrt{4 - \kappa^2\pi^2} + (8\kappa M^2 - 2\kappa^3 M^2 \pi^2)\sqrt{(1-\pi)(1+\pi)}}{\kappa^2 \pi^2 \sqrt{4 - \kappa^2 \pi^2}}$$

$N_{High}$ represents the conservative bound.

### E. Narrowing the Bounds

By providing a strict lower bound on the power of the LATE estimator, the method presented above provides a disciplined and reliable means of performing a conservative power analysis for the LATE. However, there is a tradeoff between conservatism and efficiency. If the lower bound is too conservative, it will lead to underestimation of the power and hence overestimation of the MDES and sample size required. This could then result in sub-optimal outcomes, such as a study being over-funded to achieve the conservative sample size or perhaps not funded at all if the sample size requirements exceed financial resources available. As a result, it would be useful to narrow the bounds on the power formula where possible.

As described above, the power formula bounds were based on the bounds of $Cov(\nu_1, \nu_3)$. Thus, it is possible to narrow the power bounds, specifically to increase the lower bound, if it is the case that $Cov(\nu_1, \nu_3) \geq 0$. Assuming without loss of generality that $\tau > 0$, it can be shown that this condition is met when $\bar{Y}_{NT} \leq \bar{Y}_C \leq \bar{Y}_{AT}$, where $\bar{Y}_C$, $\bar{Y}_{NT}$, and $\bar{Y}_{AT}$ denote the expected observed outcome value for compliers, never-takers, and always-takers. This will be called the "ordered-means assumption."

**Lemma 6:** *Assuming without loss of generality that* $\tau > 0$*,*

$$Cov(E[D|X], E[Y|X]) \geq 0 \Longrightarrow Cov(\nu_1, \nu_3) \geq 0$$

$$\therefore \ \bar{Y}_{NT} \leq \bar{Y}_C \leq \bar{Y}_{AT} \Longrightarrow Cov(\nu_1, \nu_3) \geq 0$$

In the case of one-sided noncompliance, the missing principal stratum will not factor into $Cov(E[D|X], E[Y|X])$. This simplifies the ordered-means assumption to $\bar{Y}_{NT} \leq \bar{Y}_C$ or $\bar{Y}_C \leq \bar{Y}_{AT}$, depending upon the direction of noncompliance. It should also be noted that, in the case where noncompliance is almost one-sided (i.e. very few never-takers or very few always-takers), the sparse principal stratum will have only a negligible impact on $Cov(E[D|X], E[Y|X])$. Thus, as a practical matter, the ordered-means assumption can be simplified to $\bar{Y}_{NT} \leq \bar{Y}_C$ or $\bar{Y}_C \leq \bar{Y}_{AT}$ as long as the sparse principal stratum is deemed sufficiently small.

As a result, if a researcher believes that the ordered-means assumption is reasonable, then s/he can assume that $Cov(\nu_1, \nu_3)$ is non-negative and hence set $Cov(\nu_1, \nu_3) = 0$ for the lower bound of the power calculation:

**Theorem 3:** *Given the ordered-means assumption, the lower bound and hence conservative power formula becomes:*

$$Power^* = \Phi\left(-c^* + \sqrt{\frac{0.5N\kappa^2\pi^2}{2-\kappa^2(\pi^2-0.5)}}\right) + \Phi\left(-c^* - \sqrt{\frac{0.5N\kappa^2\pi^2}{2-\kappa^2(\pi^2-0.5)}}\right)$$

By the same process described earlier, it is possible to solve for $\kappa$ and $N$ to derive an alternative conservative MDES and required sample size, which leads to the following:

$$\kappa^* = \frac{2M}{\sqrt{2M^2\pi^2 - M^2 + N\pi^2}}$$

$$N^* = M^2\left(\frac{\frac{4}{\kappa^2}+1}{\pi^2} - 2\right)$$

where $M = \Phi^{-1}(1-\frac{\alpha}{2}) + \Phi^{-1}(1-\beta)$

When would the ordered-means assumption be reasonable? Roughly speaking, there are two factors to consider when assessing the plausibility of this assumption. The first relates to effect heterogeneity. Specifically, it should be the case that always-takers (never-takers) select into (out of) the treatment because treatment uptake for them is associated with effects that are larger (smaller) than the average treatment effect for the compliers, or at least similarly sized. For instance, in the case of a positive and beneficial treatment, we must expect the noncomplying study subjects to be sufficiently rational that they are selecting into (out of) the treatment in anticipation of a particularly good (bad) effect on their outcome. Alternatively, selection into and out of the treatment could also be made for arbitrary reasons that are uncorrelated with effect size. The second factor relates to baseline outcome levels in the absence of the treatment. Specifically, we must expect that always-takers (never-takers) do not have baseline outcome levels that are particularly low (high) compared to that of the compliers.

Precisely when the assumption should be expected to hold, in light of the two factors described above, will certainly be context dependent. However, specific research design steps can be taken to increase its plausibility. First, studies can often be designed so as to exclude one type of noncompliance. Indeed, many experiments are designed to prevent those not assigned to the treatment from accessing it and thus ensuring the absence of always-takers. By achieving one-sided noncompliance, the analyst need only consider two principal strata, rather than three. The well-known National JTPA Study is one such example (Bloom et al., 1997). In this experimental study, subjects were randomly assigned such that they were either

given an offer to enroll in a job training program (assigned to treatment) or excluded from participating in the training for an 18-month period (assigned to control). However, many subjects given access to the job training program decided not to receive the training, resulting in a large chunk of never-takers. While there were some enterprising individuals who gained access to the job training in spite of not being assigned to it, their numbers were so small that there was virtually one-sided noncompliance.

Another research design step that can be taken is to impose reasonable restrictions on the study population of interest to ensure more similar baseline outcome levels across the principal strata. Again, the JTPA experiment is illustrative, as eligibility to participate in the experiment was restricted to those with economic disadvantages and barriers to employment. Had such restrictions not been made, the study may have included employed and/or higher income professionals who likely would have opted out of the job training program regardless of treatment assignment, thus boosting the baseline economic outcome levels of the never-takers.

In fact, as shown in the SM (Appendix C), the final results from the JTPA experiment were consistent with the ordered-means assumption in terms of an outcome variable that measured the participants' earnings in the 30-month period following their random assignment. Appendix C also presents the results of two other studies that were consistent with the ordered-means assumption. One is a vote-canvassing field experiment (Gerber and Green, 2000). The other is a fuzzy regression discontinuity design on the effect of naturalization on political integration (Hainmueller et al., 2016). That the ordered-means assumption was met in all three of these studies, which involved distinct study designs and research topics, demonstrates the plausibility of this assumption across various research domains.

## V. Power Plots

This section provides graphical illustrations of the relationships between power and the investigation parameters in the LATE power analysis method described above. Figures 1-4 each contain a set of four plots. For each set, the $x$-axis plots a particular investigation parameter while the $y$-axis plots either power or another investigation parameter. In addition, the compliance rate is varied across the four plots in each set. Every plot contains three curves. The two solid curves correspond to the lower and upper bounds. The dashed line corresponds to the alternative lower bound given the ordered-means assumption, as described in subsection IV.E. The type-I error tolerance ($\alpha$) is held constant at 0.05 for all plots.

Figure 1 plots the test power against the effect size ($\kappa$). The sample size ($N$) is set to 1000

Figure 1: Power by Effect Size (with Varying Compliance Rate)

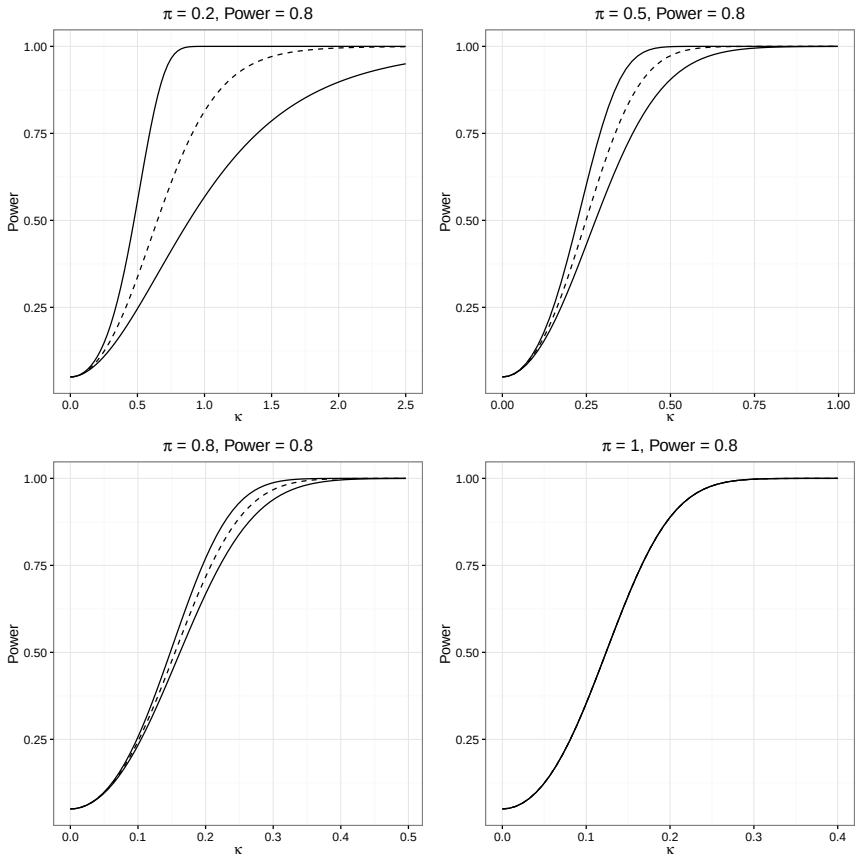


for all four plots in this figure, while the compliance rate ($\pi$) varies between 0.2, 0.5, 0.8, and 1. In addition to showing the general rate at which power increases as the effect size increases, the figure illustrates the importance of the compliance rate in shaping the relationship between power and the effect size. First, it shows that the rate at which power increases with the effect size is faster for higher compliance rates. For instance, for the alternative lower bound (the dashed line), given $\pi = 0.2$, a power of 0.8 is achieved only when $\kappa \approx 1$. In contrast, this power is achieved when $\kappa \approx 0.25$ given $\pi = 0.8$. In addition, the figure also shows that the bounds converge as $\pi \to 1$. This makes intuitive sense: when $\pi = 1$ (i.e. full compliance), the LATE becomes the ATE, which does not require bounds to characterize the power.

Figure 2 plots the test power against the sample size ($N$), with $\kappa$ held constant and $\pi$ varying across plot as previously. As in Figure 1, Figure 2 shows a similar impact of the compliance rate on the relationship between the sample size and power. Figures 3 and 4 plot the relationship between effect size and sample size, given a prespecified level of desired test power (held

Figure 2: Power by Sample Size (with Varying Compliance Rate)

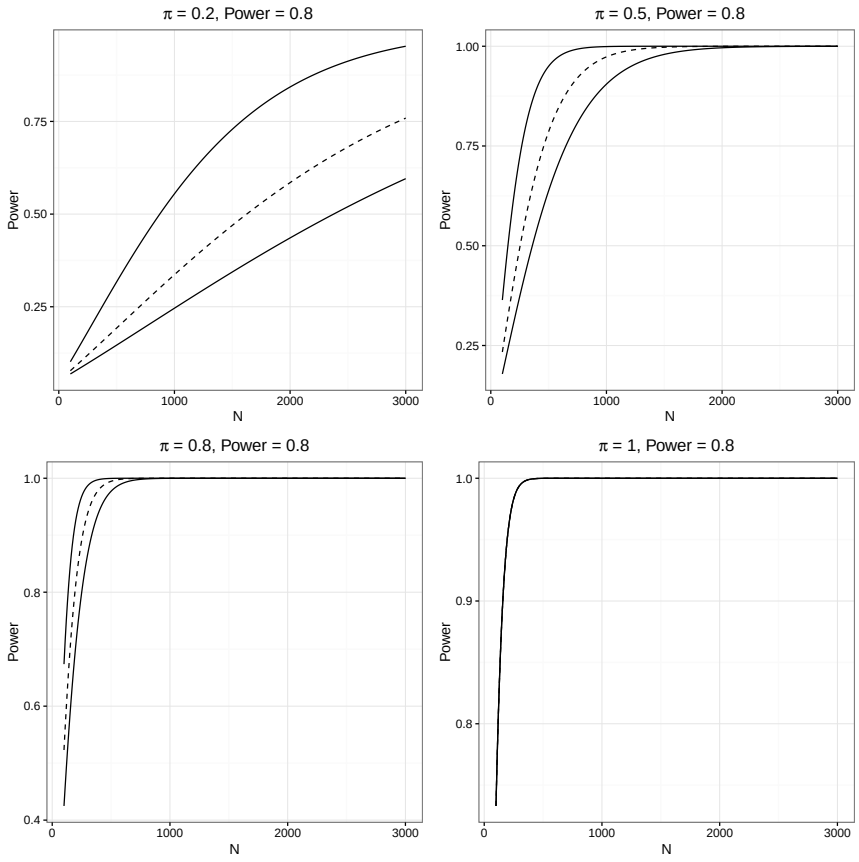


constant at 0.8). Again, the compliance rate varies across plots, illustrating the importance of the compliance rate in determining the relationship between the other parameters and the distance between the bounds.

Finally, Figure 5 compares simulated power curves to the analytic bounds, where power is plotted as a function of $\kappa$. As in Figures 1-4, the solid black lines denote the analytic upper and lower bounds, while the dashed black line denotes the alternative lower bound given the ordered-means assumption. The colored lines denote power curves that were simulated by specifying the full set of investigation and distribution parameters, and hence simulating all three principal strata, as described in section III. For all of the curves (analytic and simulated), the following parameters are fixed: $\pi = 0.2$, $N = 1000$, $\alpha = 0.05$. In addition, for the simulated power curves, the following seven of the nine distribution parameters are fixed: $E[Y_0^C] = 0$, $SD(Y_0^C) = 8$, $SD(Y_1^C) = 8$, $\dot{\pi}^{NT} = 0.4$, $SD(Y_0^{NT}) = 12$, $\dot{\pi}^{AT} = 0.4$, $SD(Y_1^{AT}) = 4$. In contrast, the final two distribution parameters, $E[Y_0^{NT}]$ and $E[Y_1^{AT}]$, vary across the five

Figure 3: MDES by Sample Size (with Varying Compliance Rate)

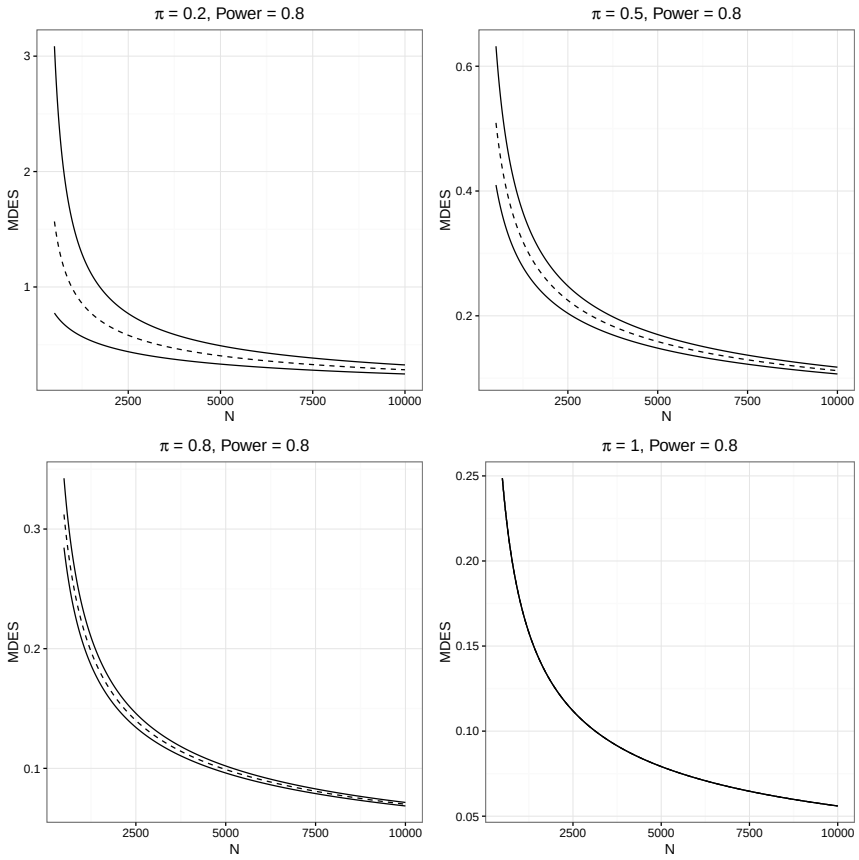


different simulation specifications shown in different colors. The five sets of values of $E[Y_0^{NT}]$ and $E[Y_1^{AT}]$, starting with the first specification, are as follows: $(-20, 20)$, $(-10, 10)$, $(-3, 3)$, $(10, -10)$, and $(20, -20)$.

Figure 5 provides a simple demonstration of the performance of the analytic bounds across these different simulation specifications. As can be seen, the simulated curves fall within the analytic bounds denoted by the solid black lines.[7] Furthermore, the alternative lower bound (the dashed black line) also bounds the appropriate simulated curves. For specifications 1 and 2, the ordered-means assumption is met at all values of $\kappa$, and hence the alternative lower bound applies. Accordingly, the curves for these specifications all lie above the alternative lower bound. In contrast, the ordered-means assumption is violated by specifications 4 and 5

[7]It appears, however, that the simulated power curves fall slightly outside the bounds at low levels of power. This is not necessarily surprising given that the power formula relies upon asymptotic normality, the LATE estimator is likely to exhibit somewhat irregular tail behavior in finite samples, and the LATE estimator's variance must itself be estimated.

Figure 4: Sample Size by Effect Size (with Varying Compliance Rate)

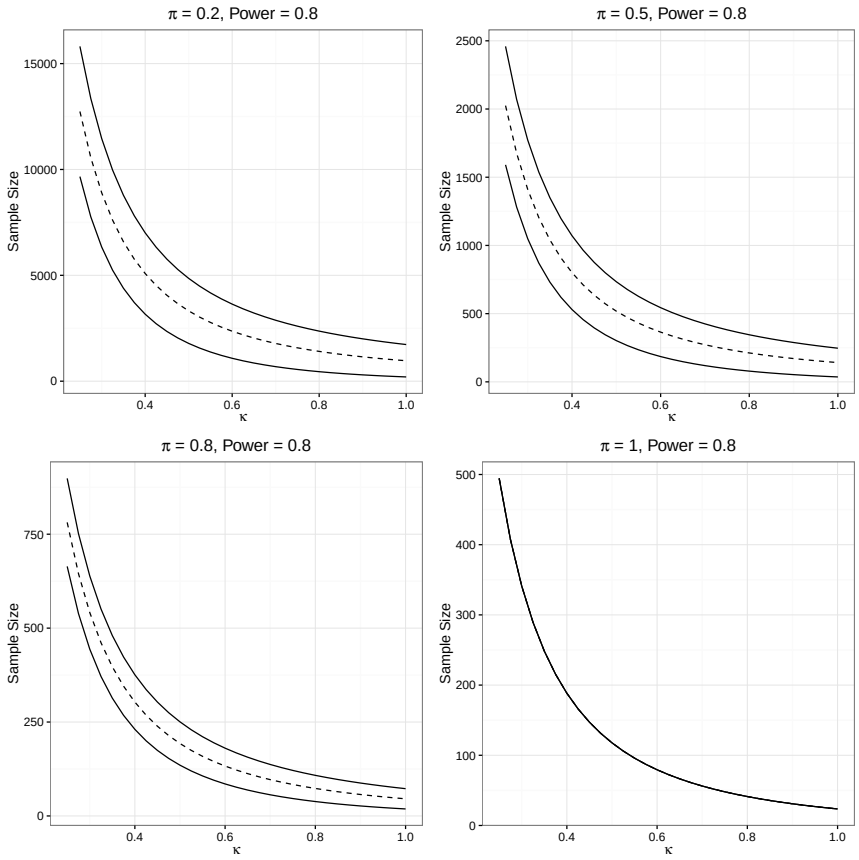


at all values of $\kappa$. Thus, it is no surprise that the curves for these specifications lie below the alternative lower bound.

## VI. Relaxing the Assumption that $P(Z = 1) = 0.5$

It some situations, the researcher may have reason to not put an equal probability on assignment to the treatment and control conditions. For instance, treatment assignment/encouragement may be costly. For such cases, it will be useful to relax the assumption that $P(Z = 1) = 0.5$.

### A. Results with $P(Z = 1) = p_z$

In the IV-LATE literature, the simplifying assumption is often made that $E[\nu_2^2|Z] = \sigma_2^2$. While there may be few cases in which this assumption is likely to hold exactly, it is often sufficiently reasonable such that it does not substantially affect statistical inference. In the method of

Figure 5: Simulated Power vs. Analytic Bounds

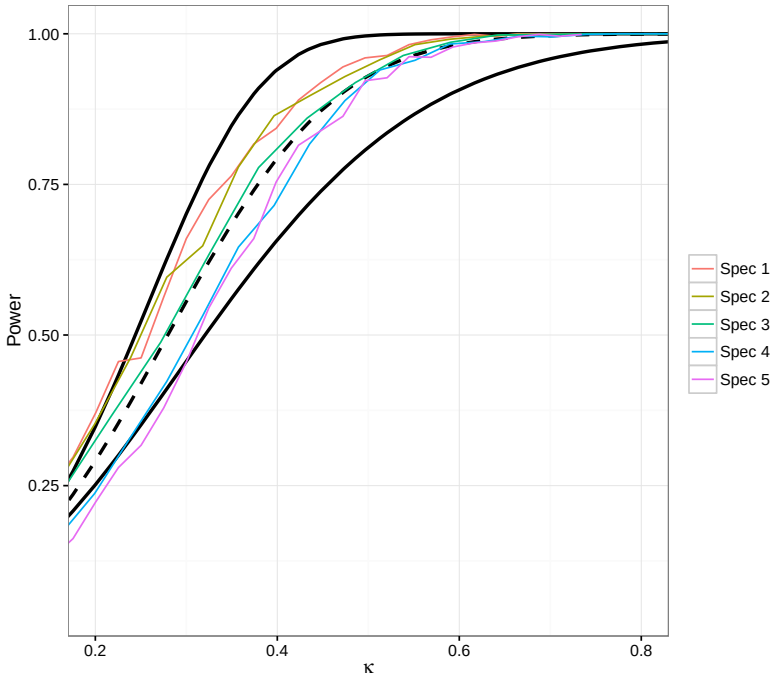


power analysis presented in this paper, the assumption that $E[\nu_2^2|Z] = \sigma_2^2$ can be useful. Specifically, if that assumption is made, then regardless of the value of $P(Z = 1)$, it is the case that:

$$\widehat{V_N^{LATE}} = \frac{E[\nu_2^2\{Z - E[Z]\}^2]}{NCov^2(D, Z)} = \frac{\sigma_2^2 Var(Z)}{NCov^2(D, Z)}$$

As a result, the assumption that $P(Z = 1) = 0.5$ is no longer necessary, as it was in Lemma 1. Lemmas 2-6 and Theorems 1-3 then proceed as before with two exceptions. First, whereas previously $Var(Z) = 0.25$, now $Var(Z) = p_z(1 - p_z)$, where $p_z = P(Z = 1)$. Hence, all terms in the power formula that include the variance of $Z$ are modified accordingly. Second, because the simplifying assumption that $E[\nu_2^2|Z] = \sigma_2^2$ is not a conservative one, it will be useful to induce conservatism elsewhere. In order to do this, we can set $\sigma_1^2 = 0.25$. This is conservative because $D$ is binary, which means its maximum marginal variance is 0.25, a value that its conditional variance ($\sigma_1^2$) cannot exceed. The following bounds on $\frac{\tau}{\sqrt{\widehat{V_N^{LATE}}}}$ then follow.

**Theorem 4:** *Given* $E[\nu_2^2|Z] = \sigma_2^2$, $\sigma_1^2 = 0.25$, *and any value* $P(Z = 1) = p_z$*:*

$$\frac{\tau}{\sqrt{\widehat{V_N^{LATE}}}}\ relaxed\ bounds: \quad \sqrt{\frac{\kappa^2(1 + p_z(1-p_z)(\frac{\kappa^2\pi^2}{1-p_z(1-p_z)\kappa^2\pi^2}))p_z(1-p_z)N\pi^2}{1 + \frac{0.25\kappa^2}{1-p_z(1-p_z)\kappa^2\pi^2} \pm \frac{\kappa}{\sqrt{1-p_z(1-p_z)\kappa^2\pi^2}}}}$$

As before, to derive bounds on the power, the bounds above can simply be plugged into the power formula $\Phi\left(-c^* + \frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}\right) + \Phi\left(-c^* - \frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}\right) = 1 - \beta$. Again, $\kappa$ can also be isolated such that the MDES can be computed as a function of the other parameters. Continuing to assume without loss of generality that $\tau > 0$, the resulting bounds on the MDES are as follows:

$$\kappa_{Low} = \sqrt{\frac{4\pi^2 J + 16\pi^2 p_z(1-p_z) - 8\sqrt{J}\pi + 4}{J^2\pi^4 + 8\pi^4 J p_z(1-p_z) - 2\pi^2 J + 16\pi^4(p_z(1-p_z))^2 + 8\pi^2 p_z(1-p_z) + 1}}$$

$$\kappa_{High} = \sqrt{\frac{4\pi^2 J + 16\pi^2 p_z(1-p_z) + 8\sqrt{J}\pi + 4}{J^2\pi^4 + 8\pi^4 J p_z(1-p_z) - 2\pi^2 J + 16\pi^4(p_z(1-p_z))^2 + 8\pi^2 p_z(1-p_z) + 1}}$$

where

$$M = \Phi^{-1}(1 - \frac{\alpha}{2}) + \Phi^{-1}(1-\beta)$$

and

$$J = \frac{4N p_z(1-p_z)}{M^2}$$

$\kappa_{High}$ represents the conservative bound.

$N$ can also be isolated in order to solve for the required sample size. Continuing to assume without loss of generality that $\tau > 0$, the resulting bounds on the sample size are as follows:

$$N_{Low} = \frac{-4\kappa\sqrt{1 - \kappa^2\pi^2 p_z(1-p_z)}\, W + \kappa^2 W + 4W - 4\kappa^2\pi^2 p_z(1-p_z)W}{\kappa^2\pi^2}$$

$$N_{High} = \frac{4\kappa\sqrt{1 - \kappa^2\pi^2 p_z(1-p_z)}\, W + \kappa^2 W + 4W - 4\kappa^2\pi^2 p_z(1-p_z)W}{\kappa^2\pi^2}$$

where

$$W = \frac{M^2}{4p_z(1-p_z)}$$

$N_{High}$ represents the conservative bound.

### B. Narrowing the Bounds while Relaxing Assumption that $P(Z = 1) = 0.5$

As before, the lower bound of the power can be increased given the ordered-means assumption ($\bar{Y}_{NT} \leq \bar{Y}_C \leq \bar{Y}_{AT}$). The result is the following alternative lower bound.

**Theorem 5:** *Given* $E[\nu_2^2|Z] = \sigma_2^2$, $\sigma_1^2 = 0.25$, *any value* $P(Z = 1) = p_z$, *and the ordered-means assumption, the lower bound and hence conservative power formula becomes:*

$$Power^* = \Phi\left(-c^* + \sqrt{\frac{\kappa^2(1 + p_z(1-p_z)(\frac{\kappa^2\pi^2}{1-p_z(1-p_z)\kappa^2\pi^2}))p_z(1-p_z)N\pi^2}{1 + \frac{0.25\kappa^2}{1-p_z(1-p_z)\kappa^2\pi^2}}}\right)$$

$$+ \Phi\left(-c^* - \sqrt{\frac{\kappa^2(1 + p_z(1-p_z)(\frac{\kappa^2\pi^2}{1-p_z(1-p_z)\kappa^2\pi^2}))p_z(1-p_z)N\pi^2}{1 + \frac{0.25\kappa^2}{1-p_z(1-p_z)\kappa^2\pi^2}}}\right)$$

Solving for $\kappa$ and $N$ to derive an alternative conservative MDES and required sample size leads to the following:

$$\kappa^* = \frac{2M}{\sqrt{4M^2\pi^2 p_z(1-p_z) - M^2 + 4N\pi^2 p_z(1-p_z)}}$$

$$N^* = \frac{\kappa^2 M^2 + 4M^2 - 4\kappa^2 M^2\pi^2 p_z(1-p_z)}{4\kappa^2\pi^2 p_z(1-p_z)}$$

where $M = \Phi^{-1}(1 - \frac{\alpha}{2}) + \Phi^{-1}(1-\beta)$

Appendix D in the SM presents results comparing simulated power curves to the analytic bounds given $P(Z = 1) = 0.25$, similar to the results shown earlier in Figure 5. Appendix D demonstrates that the analytic bounds derived for the general case where $P(Z = 1) = p_z$ perform as well as the bounds derived for the special case of $P(Z = 1) = 0.5$.

## VII. The Method in Context

To illustrate how the method of LATE power analysis presented in this paper could be used, we can implement it in the context of the National JTPA Study (Bloom et al., 1997). As described earlier, subjects were randomly assigned such that they were either allowed to enroll in a job training program (assigned to treatment) or excluded from the training for an 18-month period (assigned to control). However, many subjects exhibited noncompliance: many assigned to the treatment decided not to enroll in the training program, while a few assigned to the control gained access to the job training program. A number of economic outcomes were measured, and the outcome variable of interest here is the study subjects' earnings in the 30-month period following their random assignment.

For, the purposes of this illustration, we will employ two different values of $\pi$. This first is its observed estimated value in the JTPA data,[8] which is 0.63. This is, of course, not necessarily something the researcher would know in advance, but it provides a useful point of reference. The second value for $\pi$ will be 0.4, which we may view as a researcher's conservative guess prior to the actual study. We will fix $p_z$ at its observed value of 0.67, since this is a value over which the researcher has control, and set $\beta$ and $\alpha$ at their conventional levels of 0.2 and 0.05, respectively. We can then specify a range of effect sizes ($\kappa$'s) to determine the bounds on the sample size required under these specifications. Furthermore, because we know in retrospect the variance of the outcome (earnings), we can map the $\kappa$ values to absolute effect values ($\tau$'s).

The results given $\pi = 0.63$ are shown in Table 1, with the conservative $N$ corresponding to the power lower bound, and the alternative $N$ corresponding to the alternative power lower bound given the ordered-means assumption. For instance, given a desired LATE effect size of 0.1, the conservative sample size recommendation would be approximately $10,000$ observations to achieve a test of power of 0.8, while the alternative sample size given the order-means assumption would be approximately $9,000$ observations. The actual LATE effect size estimate in the pooled adult sample was 0.11.[9] Thus, it is no surprise that given the actual sample size of $11,204$ adult participants and the fact that the ordered-means assumption was ultimately met in this study, the estimated LATE is strongly statistically significant ($p < 0.001$).

Table 1: LATE Power Analysis, Given $\pi = 0.63$

| $\kappa$ | $\tau$ | Conservative $N$ from LATE Power Analysis | Alternative $N$ from LATE Power Analysis |
|---|---|---|---|
| 0.05 | 838.35 | 37824 | 36024 |
| 0.10 | 1676.70 | 9917 | 9017 |
| 0.15 | 2515.06 | 4615 | 4016 |
| 0.20 | 3353.41 | 2715 | 2265 |
| 0.25 | 4191.76 | 1814 | 1455 |
| 0.30 | 5030.11 | 1314 | 1015 |
| 0.35 | 5868.47 | 1005 | 750 |
| 0.40 | 6706.82 | 801 | 577 |
| 0.45 | 7545.17 | 657 | 459 |
| 0.50 | 8383.52 | 553 | 375 |

Table 2 displays the results given $\pi = 0.4$. As shown, an increase in the amount of noncompliance leads to a disproportionately large increase in the sample size requirements.

[8]The dataset used here is the tabulation of the JTPA study data by Abadie et al. (2002). The data correspond to adult participants in the JTPA experiment for whom 30-month earnings were measured.

[9]The estimate of the local average treatment effect of the training program on earnings is $\$1,849$. This divided by the standard deviation of earnings in the sample, $\$16,767$, yields an effect size of 0.11.

While noncompliance is assumed to increase by a factor of about 1.6, the required sample size given any particular $\kappa$ under both bounds increases by a factor of about 2.5. As these results show, had the compliance rate $\pi$ actually been 0.4, it is likely that the JTPA study would have failed to find a statistically significant local effect of the training program on earnings, even in the pooled adult sample. The method of LATE power analysis presented in this paper is designed to alert researchers to such possibilities of under-powered designs before studies are launched without requiring researchers to make the collection of assumptions required by alternative approaches to LATE power analysis.

Table 2: LATE Power Analysis, Given $\pi = 0.4$

| $\kappa$ | $\tau$ | Conservative $N$ from LATE Power Analysis | Alternative $N$ from LATE Power Analysis |
|---|---|---|---|
| 0.05 | 838.35 | 92983 | 88558 |
| 0.10 | 1676.70 | 24387 | 22175 |
| 0.15 | 2515.06 | 11357 | 9882 |
| 0.20 | 3353.41 | 6685 | 5579 |
| 0.25 | 4191.76 | 4472 | 3588 |
| 0.30 | 5030.11 | 3242 | 2506 |
| 0.35 | 5868.47 | 2485 | 1854 |
| 0.40 | 6706.82 | 1982 | 1430 |
| 0.45 | 7545.17 | 1630 | 1140 |
| 0.50 | 8383.52 | 1373 | 933 |

## VIII. Conclusion

This paper proposes a new approach to power analysis in the LATE context that makes four important contributions. First, in contrast to existing approaches, it does not involve distributional assumptions about the various principal strata. Second, it operationalizes the effect size of interest in meaningful relative terms, using the most natural reference variance in the LATE context, which is the variance of the full distribution of the observed outcome. Third, and most importantly, it provides a strict lower bound on the power calculation while removing the need to specify or make assumptions about variance components or distributional heterogeneity across the principal strata. Fourth, it shows how additional assumptions can be made to raise the lower bound to better balance conservatism with efficiency.

By providing bounds on the power calculation that are free of distributional assumptions, this paper introduces a reliable and disciplined way of computing power conservatively without the inefficiencies of other approaches (e.g. setting arbitrarily high variances) that can lead to

excessively conservative calculations. The result is a generalized approach to power analysis in the LATE context that is simultaneously conservative, disciplined, and simple to implement. This is an advance over preexisting power analysis methods in the LATE context, which require making explicit guesses or implicit assumptions about reasonable treatment effect magnitudes, multiple principal strata mean and variance components, and the proportions of the population belonging to each principal stratum.

## References


Abadie, A., Angrist, J. D., and Imbens, G. W. (2002). Instrumental variables estimates of the effect of subsidized training on the quantiles of trainee earnings. *Econometrica*, 70(1):91–117.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455.

Baiocchi, M., Cheng, J., and Small, D. S. (2014). Tutorial in biostatistics: Instrumental variable methods for causal inference. *Statistics in Medicine*, 33(13):2297–2340.

Bloom, H. S. (1995). Minimum detectable effects: A simple way to report the statistical power of experimental designs. *Evaluation Review*, 19(5):547–556.

Bloom, H. S. (2006). The core analytics of randomized experiments for social research. Technical report, MDRC.

Bloom, H. S., Orr, L. L., Bell, S. H., Cave, G., Doolittle, F., Lin, W., and Bos, J. M. (1997). The benefits and costs of JTPA Title II-A programs: Key findings from the national job training partnership act study. *The Journal of Human Resources*, 32:549–576.

Brion, M.-J. A., Shakhbazov, K., and Visscher, P. M. (2013). Calculating statistical power in mendelian randomization studies. *International Journal of Epidemiology*, 42(5):1497–1501.

Cohen, J. (1962). The statistical power of abnormal-social psychological research. *Journal of Abnormal and Social Psychology*, 65:145–153.

Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Duflo, E., Glennerster, R., and Kremer, M. (2007). Using randomization in development economics research: A toolkit. Technical report, Centre for Economic Policy Research.

Freeman, G., Cowling, B. J., and Schooling, C. M. (2013). Power and sample size calculations for mendelian randomization studies using one genetic instrument. *International Journal of Epidemiology*, 42(4):1157–1163.

Gerber, A. S. and Green, D. P. (2000). The effects of canvassing, telephone calls, and direct mail on voter turnout: A field experiment. *American Political Science Review*, 94(3):653–663.

Gerber, A. S. and Green, D. P. (2005). Correction to gerber and green (2000), replication of disputed findings, and reply to imai (2005). *American Political Science Review*, 99(2):301–313.

Gerber, A. S. and Green, D. P. (2012). *Field Experiments: Design, Analysis, and Interpretation*. New York: W. W. Norton & Company.

Hahn, J., Todd, P., and Van der Klaauw, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209.

Hainmueller, J., Hangartner, D., and Pietrantuono, G. (2016). Naturalization fosters the long-term political integration of immigrants. *Proceedings of the National Academy of Sciences*, 112(41):12651–12656.

Imbens, G. W. and Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475.

Imbens, G. W. and Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge: Cambridge University Press.

Jo, B. (2002). Statistical power in randomized intervention studies with noncompliance. *Psychological Methods*, 7(2):178–193.

Lipsey, M. W. (1990). *Design Sensitivity: Statistical Power for Experimental Research*. Newbury Park, CA: Sage Publications.

Pierce, B. L., Ahsan, H., and VanderWeele, T. J. (2011). Power and instrument strength requirements for mendelian randomization studies using multiple genetic variants. *International Journal of Epidemiology*, 40(3):740–752.

Tsang, R., Colley, L., and Lynd, L. D. (2009). Inadequate statistical power to detect clinically significant differences in adverse event rates in randomized controlled trials. *Journal of Clinical Epidemiology*, 62(6):609–616.

Tversky, A. and Kahneman, D. (1971). The belief in the law of small numbers. *Psychological Bulletin*, 76(2):105–110.

**Supplementary Materials**

for

# A Generalized Approach to Power Analysis for Estimation of Local Average Treatment Effects

Kirk Bansak

# Appendix A: Proofs

*Proof of Lemma 1.*

Given $P(Z=1)=0.5$:
$E[\nu_2^2\{Z-E[Z]\}^2] = E[\nu_2^2\{Z-0.5\}^2] = E[E[\nu_2^2\{Z-0.5\}^2|Z]] = p(Z=0)E[\nu_{2(Z=0)}^2\{0-0.5\}^2]+p(Z=1)E[\nu_{2(Z=1)}^2\{1-0.5\}^2] = 0.5E[\nu_{2(Z=0)}^2 0.25]+0.5E[\nu_{2(Z=1)}^2 0.25] = \{0.5E[\nu_{2(Z=0)}^2]+0.5E[\nu_{2(Z=1)}^2]\}0.25 = E[\nu_2^2]0.25$

Thus,

$$\widehat{V_N^{LATE}} = \frac{E[\nu_2^2\{Z-E[Z]\}^2]}{NCov^2(D,Z)} = \frac{E[\nu_2^2]Var(Z)}{NCov^2(D,Z)}$$

Further, this can be re-expressed as:

$$\widehat{V_N^{LATE}} = \frac{E[\nu_2^2]}{N}\frac{Var(Z)}{Cov^2(D,Z)}\frac{Var(Z)}{Var(Z)} = \frac{E[\nu_2^2]}{N}\frac{Var(Z)}{Cov(D,Z)}\frac{Var(Z)}{Cov(D,Z)}\frac{1}{Var(Z)}$$

Finally, since $\frac{Cov(D,Z)}{Var(Z)} = \pi$ and $Var(Z) = 0.25$, this further simplifies to:

$$\widehat{V_N^{LATE}} = \frac{E[\nu_2^2]}{0.25N\pi^2}$$

□

*Proof of Lemma 2.*

Note that $\nu_2 = \nu_3 - \tau\nu_1$ by equations (1)-(3) and the property that $\tau = \frac{\gamma}{\pi}$ under the LATE assumptions. Thus,

$$E[\nu_2^2] = E[(\nu_3 - \tau\nu_1)^2] = E[\nu_3^2] + \tau^2 E[\nu_1^2] - 2\tau E[\nu_1\nu_3] = \sigma_3^2 + \tau^2\sigma_1^2 - 2\tau\sigma_{13}$$

where $\sigma_1^2 = Var(\nu_1)$, $\sigma_3^2 = Var(\nu_3)$, and $\sigma_{13} = Cov(\nu_1, \nu_3)$, since $E[\nu_1] = E[\nu_3] = 0$.

Combined with Lemma 1, this results in:

$$\widehat{V_N^{LATE}} = \frac{\sigma_3^2 + \tau^2\sigma_1^2 - 2\tau\sigma_{13}}{0.25N\pi^2}$$

□

*Proof of Lemma 3.*

$\sigma_{13}$ is bounded as follows:

$$-\sigma_1\sigma_3 = -\sqrt{Var(\nu_1)Var(\nu_3)} \leq \sigma_{1,3} = Cov(\nu_1, \nu_3) \leq \sqrt{Var(\nu_1)Var(\nu_3)} = \sigma_1\sigma_3$$

Hence, assuming without loss of generality that $\tau > 0$, it is conservative to set $Cov(\nu_1, \nu_3)$ to its lower-bound value (thereby setting $-2\tau\sigma_{13}$ to its upper-bound value):

$$\sigma_3^2 + \tau^2\sigma_1^2 - 2\tau\sigma_{13} \equiv \sigma_3^2 + \tau^2\sigma_1^2 - 2\tau(-\sigma_1\sigma_3) = \sigma_3^2 + \frac{\gamma^2}{\pi^2}\sigma_1^2 + 2\frac{\gamma}{\pi}\sigma_1\sigma_3$$

As a result, this component in the denominator of $\frac{\tau}{V_N^{\widehat{LATE}}}$ is set to its largest possible value, making $\frac{\tau}{V_N^{\widehat{LATE}}}$ as small as possible. This will yield a lower bound on the power formula, thus providing a reliable conservative computation. An upper bound on the power can also be computed by setting $Cov(\nu_1, \nu_3)$ to its upper-bound value. □

*Proof of Lemma 4.*

By the law of total variance:

$$Var(Y) = E[Var(Y|Z)] + Var(E[Y|Z])$$

which, given $P(Z = 1) = 0.5$ and equation (3), simplifies as:

$$Var(Y) = E[\nu_3^2] + 0.25\gamma^2 = \sigma_3^2 + 0.25\gamma^2$$

Hence:

$$\tau = \kappa\sqrt{\sigma_3^2 + 0.25\gamma^2}$$

Finally, because we also know that $\tau = \frac{\gamma}{\pi}$, then $\frac{\gamma}{\pi} = \kappa\sqrt{\sigma_3^2 + 0.25\gamma^2}$, which can be solved for $\gamma$ as:

$$\gamma = \sqrt{\frac{\kappa^2\sigma_3^2\pi^2}{1 - 0.25\kappa^2\pi^2}}$$

□

*Proof of Theorem 1.*

For convenience, work with the square of $\frac{\tau}{\sqrt{V_N^{\widehat{LATE}}}}$. Given $P(Z=1)=0.5$, definition 1, and lemmas 1-4:

$$\begin{aligned}
&\frac{\tau^2}{V_N^{\widehat{LATE}}} \\
&= \frac{\kappa^2(\sigma_3^2 + 0.25\gamma^2)0.25N\pi^2}{\sigma_3^2 + \frac{\gamma^2}{\pi^2}\sigma_1^2 \pm 2\frac{\gamma}{\pi}\sigma_1\sigma_3} \\
&= \frac{\kappa^2(\sigma_3^2 + 0.25(\frac{\kappa^2\pi^2\sigma_3^2}{1-0.25\kappa^2\pi^2}))0.25N\pi^2}{\sigma_3^2 + \frac{\kappa^2\pi^2\sigma_1^2\sigma_3^2}{\pi^2(1-0.25\kappa^2\pi^2)} \pm \frac{2\kappa\pi\sigma_1\sigma_3\sigma_3}{\pi\sqrt{1-0.25\kappa^2\pi^2}}} \\
&= \frac{\kappa^2(1 + 0.25(\frac{\kappa^2\pi^2}{1-0.25\kappa^2\pi^2}))0.25N\pi^2}{1 + \frac{\kappa^2\sigma_1^2}{1-0.25\kappa^2\pi^2} \pm \frac{2\kappa\sigma_1}{\sqrt{1-0.25\kappa^2\pi^2}}}
\end{aligned}$$

Notably, $\sigma_3^2$ can be factored out of both the numerator and denominator and hence canceled out, thus eliminating a variance term. $\square$

*Proof of Lemma 5.*

Consider that, given P(Z=1)=0.5:

$$\sigma_1^2 = E[Var(D|Z)] = 0.5Var(D|Z=0)+0.5Var(D|Z=1) = 0.5[Var(D|Z=0)+Var(D|Z=1)]$$

Thus, $\sigma_1^2$ is maximized by maximizing $Var(D|Z=0) + Var(D|Z=1)$.

Further, consider that $D$ is binary with $E[D|Z=1] = E[D|Z=0] + \pi$. As a result, $Var(D|Z=0)+Var(D|Z=1)$ is maximized when $E[D|Z=1]$ and $E[D|Z=0]$ are equidistant from 0.5. This occurs when $E[D|Z=0] = 0.5 - \frac{\pi}{2}$ and $E[D|Z=0] = 0.5 + \frac{\pi}{2}$. The result is the following upper-bound value:

$$\begin{aligned}\sigma_1^2 = E[Var(D|Z)] &= 0.5Var(D|Z=0) + 0.5Var(D|Z=1) \\ &= 0.5(0.5 - \frac{\pi}{2})(0.5 + \frac{\pi}{2}) + 0.5(0.5 + \frac{\pi}{2})(0.5 - \frac{\pi}{2}) \\ &= (0.5 - \frac{\pi}{2})(0.5 + \frac{\pi}{2})\end{aligned}$$

□

*Proof of Theorem 2.*

Simply substitute $(0.5 - \frac{\pi}{2})(0.5 + \frac{\pi}{2})$ in for $\sigma_1^2$ in the result in Theorem 1. $\square$

*Proof of Lemma 6.*

First, consider the following re-expressions of equations (1) and (3):

$$\nu_1 = D - E[D] - \pi Z + \pi E[Z]$$
$$\nu_3 = Y - E[Y] - \gamma Z + \gamma E[Z]$$

Hence

$$\begin{aligned} Cov(\nu_1, \nu_3) &= Cov(D - \pi Z, Y - \gamma Z) \\ &= Cov(D,Y) - \pi Cov(Z,Y) - \gamma Cov(D,Z) + \pi\gamma Var(Z) \\ &= Cov(D,Y) - \pi Cov(Z,Y) - \gamma Cov(D,Z) + \frac{Cov(D,Z)}{Var(Z)}\gamma Var(Z) \\ &= Cov(D,Y) - \pi Cov(Z,Y) \end{aligned}$$

It is further possible to decompose $Cov(D,Y)$ using the law of total covariance, conditioning on principal strata. Let $X$ denote the principal strata, making the usual assumption of no defiers. Continue to assume that $P(Z=1) = 0.5$, and WLOG that $\tau > 0$. Let $\bar{Y}_C$, $\bar{Y}_{NT}$, and $\bar{Y}_{AT}$ denote the expected observed outcome value for compliers, never-takers, and always-takers. Then,

$$\begin{aligned} Cov(D,Y) &= E[Cov(D,Y|X)] + Cov(E[D|X], E[Y|X]) \\ &= P(Compliers)Cov(D,Y|Compliers) + Cov(E[D|X], E[Y|X]) \end{aligned}$$

since $Cov(D,Y|Always-Takers) = Cov(D,Y|Never-Takers) = 0$.

Furthermore, we can show that:

$$\begin{aligned} P(Compliers)Cov(D,Y|Compliers) &= P(Compliers)Cov(Z,Y|Compliers) \\ &= \pi\frac{Cov(Z,Y|Compliers)}{Var(Z|Compliers)}Var(Z|Compliers) \\ &= \pi\frac{Cov(D,Y|Compliers)}{Var(D|Compliers)}Var(Z) \\ &= \pi\tau Var(Z) \\ &= \pi\frac{\gamma}{\pi}Var(Z) \\ &= \gamma Var(Z) \\ &= \frac{Cov(Z,Y)}{Var(Z)}Var(Z) \\ &= Cov(Z,Y) \end{aligned}$$

This means that:

$$\begin{aligned} Cov(\nu_1, \nu_3) &= Cov(D, Y) - \pi Cov(Z, Y) \\ &= Cov(Z, Y) + Cov(E[D|X], E[Y|X]) - \pi Cov(Z, Y) \\ &= (1 - \pi) Cov(Z, Y) + Cov(E[D|X], E[Y|X]) \end{aligned}$$

The first term $(1 - \pi)Cov(Z, Y)$ is non-negative by the assumption (WLOG) that $\tau > 0$. Hence, $Cov(\nu_1, \nu_3)$ is guaranteed to be non-negative if $Cov(E[D|X], E[Y|X])$ is non-negative, which is true when $\bar{Y}_{NT} \leq \bar{Y}_C \leq \bar{Y}_{AT}$. $\square$

*Proof of Theorem 3.*

By Lemma 6, $\bar{Y}_{NT} \leq \bar{Y}_C \leq \bar{Y}_{AT} \implies Cov(\nu_1, \nu_3) \geq 0$. The value that produces the most conservative (i.e. lowest) power calculation is $Cov(\nu_1, \nu_3) = 0$. Given $Cov(\nu_1, \nu_3) = 0$,

$$\begin{aligned}
\frac{\tau^2}{\widehat{V_N^{LATE}}} &= \frac{\kappa^2(\sigma_3^2 + 0.25\gamma^2)0.25N\pi^2}{\sigma_3^2 + \frac{\gamma^2}{\pi^2}\sigma_1^2} \\
&= \frac{\kappa^2(\sigma_3^2 + 0.25(\frac{\kappa^2\sigma_3^2\pi^2}{1-0.25\kappa^2\pi^2}))0.25N\pi^2}{\sigma_3^2 + \frac{\left(\frac{\kappa^2\sigma_3^2\pi^2}{1-0.25\kappa^2\pi^2}\right)}{\pi^2}\sigma_1^2} \\
&= \frac{\sigma_3^2\kappa^2(1 + 0.25(\frac{\kappa^2\pi^2}{1-0.25\kappa^2\pi^2}))0.25N\pi^2}{\sigma_3^2\left(1 + \left(\frac{\kappa^2}{1-0.25\pi^2\kappa^2}\right)\sigma_1^2\right)} \\
&= \frac{\kappa^2(1 + 0.25(\frac{\kappa^2\pi^2}{1-0.25\kappa^2\pi^2}))0.25N\pi^2}{\left(1 + \left(\frac{\kappa^2}{1-0.25\kappa^2\pi^2}\right)\sigma_1^2\right)} \\
&= \frac{N\kappa^2\pi^2}{4 - \kappa^2(\pi^2 - 4\sigma_1^2)}
\end{aligned}$$

Finally, since $D$ is binary, $\sigma_1^2$ can be simplified as a function of $\pi$, with $\sigma_1^2 = (0.5-\frac{\pi}{2})(0.5+\frac{\pi}{2})$ being the most conservative possibility by Lemma 5.

The result, then, is that:

$$\frac{\tau}{\sqrt{\widehat{V_N^{LATE}}}} = \sqrt{\frac{N\kappa^2\pi^2}{4 - \kappa^2(\pi^2 - 4(0.5 - \frac{\pi}{2})(0.5 + \frac{\pi}{2}))}}$$

which simplifies to:

$$\frac{\tau}{\sqrt{\widehat{V_N^{LATE}}}} = \sqrt{\frac{0.5N\kappa^2\pi^2}{2 - \kappa^2(\pi^2 - 0.5)}}$$

Hence, the lower bound for the power becomes:

$$Power^* = \Phi\left(-c^* + \sqrt{\frac{0.5N\kappa^2\pi^2}{2 - \kappa^2(\pi^2 - 0.5)}}\right) + \Phi\left(-c^* - \sqrt{\frac{0.5N\kappa^2\pi^2}{2 - \kappa^2(\pi^2 - 0.5)}}\right)$$

□

*Proof of Theorem 4.*

The results of Lemma 1 now hold by the assumption that $E[\nu_2^2|Z] = \sigma_2^2$, rather than the now discarded assumption that $P(Z = 1) = 0.5$, but with the difference that now $Var(Z) = p_z(1 - p_z) \neq 0.25$. Thus, now $\widehat{V_N^{LATE}} = \frac{\sigma_2^2}{p_z(1-p_z)N\pi^2}$. The results of Lemmas 2 and 3 proceed as before. Making use of the results of Lemma 4, but allowing $P(Z = 1) = p_z$ for any value of $p_z$, then:

$$Var(Y) = E[\nu_3^2] + Var(Z)\gamma^2 = \sigma_3^2 + Var(Z)\gamma^2$$

Hence:

$$\tau = \kappa\sqrt{\sigma_3^2 + Var(Z)\gamma^2} = \kappa\sqrt{\sigma_3^2 + p_z(1 - p_z)\gamma^2}$$

And:

$$\gamma = \sqrt{\frac{\kappa^2\sigma_3^2\pi^2}{1 - Var(Z)\kappa^2\pi^2}} = \sqrt{\frac{\kappa^2\sigma_3^2\pi^2}{1 - p_z(1 - p_z)\kappa^2\pi^2}}$$

Theorem 1 then proceeds as before but now using the new results described here. Finally, 0.25 is substituted in for $\sigma_1^2$. □

*Proof of Theorem 5.*

The proof follows from the results of Lemma 6 applied to Theorem 4. □

# Appendix B: LATE Power Simulations

This appendix presents the results of a series of simulations illustrating the power of the LATE estimator as the distributional characteristics of the principal strata are varied.

Table B1: LATE Power Simulations, Varying the Investigation Parameters

| Fixed Parameters: | Varying Parameter | Power |
|---|---|---|
| $\tau = 5$, $\pi = 0.2$, | $N =$ | |
| $E[Y_0^C] = 0$, $SD(Y_0^C) = 8$, $SD(Y_1^C) = 8$, | 1000 | **0.43** |
| $\dot{\pi}^{NT} = 0.4$, $E[Y_0^{NT}] = -3$, $SD(Y_0^{NT}) = 12$, | 2000 | **0.71** |
| $\dot{\pi}^{AT} = 0.4$, $E[Y_1^{AT}] = 3$, $SD(Y_1^{AT}) = 4$ | 4000 | **0.93** |
| $N = 1000$, $\pi = 0.2$, | $\tau =$ | |
| $E[Y_0^C] = 0$, $SD(Y_0^C) = 8$, $SD(Y_1^C) = 8$, | 5 | **0.43** |
| $\dot{\pi}^{NT} = 0.4$, $E[Y_0^{NT}] = -3$, $SD(Y_0^{NT}) = 12$, | 6 | **0.57** |
| $\dot{\pi}^{AT} = 0.4$, $E[Y_1^{AT}] = 3$, $SD(Y_1^{AT}) = 4$ | 7 | **0.68** |

In the simulations presented, $\alpha$ is held fixed at 0.05. The sample size ($N$), LATE ($\tau$), and first-stage effect ($\pi$) are specified, and samples are generated with specified proportions of compliers (which is also the first-stage effect, $\pi$), never-takers ($\dot{\pi}^{NT}$), and always-takers ($\dot{\pi}^{AT}$). For the compliers, a control condition outcome mean ($E[Y_0^C]$) is specified (the treatment condition mean is determined by the specified $\tau$), as are control condition and treatment condition outcome standard deviations ($SD(Y_0^C)$ and $SD(Y_1^C)$), with outcome values being randomly generated to be normally distributed. The outcome distributions of the never-takers and always-takers are also randomly generated to be normally distributed; for the never-takers, only a control condition mean and variance ($E[Y_0^{NT}]$ and $SD(Y_0^{NT})$) must be specified, and for the always-takers, only a treatment condition mean and variance ($E[Y_1^{AT}]$ and $SD(Y_1^{AT})$) must be specified. Finally, the treatment is randomly assigned with probability one half. However, whether a unit is actually treated, and hence obtains an observed outcome value that is the treatment (or control) condition outcome, is determined not only by treatment assignment but also the principal stratum to which that unit belongs.

A variety of specifications are presented in Tables B1 and B2, with specific parameters being varied as others are held constant to demonstrate the impact on the power of the test. For each specification, $5,000$ samples are simulated, and for each simulation, a hypothesis test of the LATE estimate is undertaken. This allows for a simulated power to be computed by calculating the proportion of the $5,000$ simulations for which it is possible to reject the null hypothesis that $\tau = 0$ given a two-tailed test with a type-I error tolerance of $\alpha = 0.05$.

Table B2: LATE Power Simulations, Varying the Distribution Parameters

| Fixed Parameters: | Varying Parameter(s) | Power |
|---|---|---|
| $N = 1000$, $\tau = 5$, $\pi = 0.2$, | $E[Y_0^C] =$ | |
| $SD(Y_0^C) = 8$, $SD(Y_1^C) = 8$, | 0 | **0.45** |
| $\dot{\pi}^{NT} = 0.4$, $E[Y_0^{NT}] = -3$, $SD(Y_0^{NT}) = 12$, | 10 | **0.35** |
| $\dot{\pi}^{AT} = 0.4$, $E[Y_1^{AT}] = 3$, $SD(Y_1^{AT}) = 4$ | 20 | **0.25** |
| $N = 1000$, $\tau = 5$, $\pi = 0.3$, | $SD(Y_0^C), SD(Y_1^C) =$ | |
| $E[Y_0^C] = 0$, | $8, 8$ | **0.76** |
| $\dot{\pi}^{NT} = 0.35$, $E[Y_0^{NT}] = -3$, $SD(Y_0^{NT}) = 12$, | $8, 16$ | **0.63** |
| $\dot{\pi}^{AT} = 0.35$, $E[Y_1^{AT}] = 3$, $SD(Y_1^{AT}) = 4$ | $16, 16$ | **0.52** |
| $N = 1000$, $\tau = 5$, $\pi = 0.2$, | $E[Y_0^{NT}], E[Y_1^{AT}] =$ | |
| $E[Y_0^C] = 0$, $SD(Y_0^C) = 8$, $SD(Y_1^C) = 8$, | $-3, 3$ | **0.44** |
| $\dot{\pi}^{NT} = 0.4$, $SD(Y_0^{NT}) = 12$, | $10, 3$ | **0.26** |
| $\dot{\pi}^{AT} = 0.4$, $SD(Y_1^{AT}) = 4$ | $10, -6$ | **0.13** |
| $N = 1000$, $\tau = 5$, $\pi = 0.2$, | $SD(Y_0^{NT}), SD(Y_1^{AT}) =$ | |
| $E[Y_0^C] = 0$, $SD(Y_0^C) = 8$, $SD(Y_1^C) = 8$, | $12, 4$ | **0.43** |
| $\dot{\pi}^{NT} = 0.4$, $E[Y_0^{NT}] = -3$, | $12, 8$ | **0.37** |
| $\dot{\pi}^{AT} = 0.4$, $E[Y_1^{AT}] = 3$ | $24, 8$ | **0.15** |
| $N = 1000$, $\tau = 5$, | $\pi$, $\dot{\pi}^{NT}, \dot{\pi}^{AT} =$ | |
| $E[Y_0^C] = 0$, $SD(Y_0^C) = 8$, $SD(Y_1^C) = 8$, | $0.3, 0.35, 0.35$ | **0.76** |
| $E[Y_0^{NT}] = -3$, $SD(Y_0^{NT}) = 12$, | $0.2, 0.4, 0.4$ | **0.45** |
| $E[Y_1^{AT}] = 3$, $SD(Y_1^{AT}) = 4$ | $0.2, 0.1, 0.7$ | **0.71** |
| | $0.2, 0.8, 0.0$ | **0.30** |

The results shown in Tables B1 and B2 demonstrate that, in addition to the investigation parameters, all nine of the distribution parameters have an impact on the power of the test.

Table B3 compares the test power of the LATE estimator with that of the ITT estimator, highlighting a phenomenon recognized in previous work by Jo (2002). Specifically, Table B3 varies the mean outcome values of the never-takers and always-takers, holding all other parameters constant. As can be seen, the two test powers change at different rates as these two principal strata means are altered. Depending upon the specification, the power of the LATE estimator may be higher or lower than the power of the ITT estimator. Table B3 demonstrates that a power test for the ITT cannot serve as a replacement for a LATE power analysis.

Table B3: LATE vs. ITT Power Simulations

| Fixed Parameters: | Varying Parameter(s) | $\text{Power}_{LATE}$ | $\text{Power}_{ITT}$ |
|---|---|---|---|
| | $E[Y_0^{NT}], E[Y_1^{AT}] =$ | | |
| $N = 1000$, $\tau = 5$, $\pi = 0.2$, | $-10, 10$ | **0.34** | **0.25** |
| $E[Y_0^C] = 0$, $SD(Y_0^C) = 8$, $SD(Y_1^C) = 8$, | $-10, 3$ | **0.38** | **0.30** |
| $\dot{\pi}^{NT} = 0.4$, $SD(Y_0^{NT}) = 12$, | $-3, 3$ | **0.44** | **0.41** |
| $\dot{\pi}^{AT} = 0.4$, $SD(Y_1^{AT}) = 4$ | $10, 3$ | **0.26** | **0.37** |
| | $10, -6$ | **0.13** | **0.30** |

Finally, Table B4 compares the test power of the LATE estimator with the power implied by scaling a standard ATE power analysis. As described by previous studies (Duflo et al., 2007; Baiocchi et al., 2014), one approach to computing the power of the LATE estimator is by performing a simple ATE power analysis but scaling the variance of the standard ATE estimator, specifically dividing it by the compliance rate squared. This implies that, for a given treatment effect, the sample size required to reach a particular power level in the full-compliance ATE context must be divided by the compliance rate squared to offer the sample size required to reach that same power level in the LATE context. For instance, if 100 observations are needed in the ATE context, then approximately $\frac{100}{0.2^2} = 2500$ observations are required in the LATE context given a compliance rate of 0.2. However, this approximate equality is only achieved given the strong assumptions that (a) the never-takers have the same mean outcome value as the untreated compliers, (b) the always-takers have the same mean outcome value as the treated compliers, and (c) all groups have the same outcome variance. If any of those assumptions are violated, the true power of the LATE estimator can diverge dramatically from the power implied by this scaled ATE power analysis.

Table B4: LATE vs. Scaled ATE Power Simulations

| Fixed Parameters: | Varying Parameter(s) | $\text{Power}_{LATE}$ | Power Implied by Scaled ATE Analysis |
|---|---|---|---|
| | $E[Y_0^{NT}], E[Y_1^{AT}], SD(Y_0^{NT}), SD(Y_1^{AT}) =$ | | |
| $N = 2500$, $\tau = 5$, $\pi = 0.2$, | $0, 5, 8, 8$ | **0.87** | · |
| $E[Y_0^C] = 0$, | $-10, 15, 8, 8$ | **0.54** | · |
| $SD(Y_0^C) = 8$, $SD(Y_1^C) = 8$, | $0, 5, 16, 16$ | **0.39** | **0.87** |
| $\dot{\pi}^{NT} = 0.4$, $\dot{\pi}^{AT} = 0.4$ | $-10, 15, 16, 16$ | **0.32** | · |
| | $15, -10, 16, 16$ | **0.19** | · |

Table B4 demonstrates this result. Specifically, a standard ATE power analysis is performed given a treatment effect of 5, a sample size of 100 (assigned equally to the treatment and control groups), and within-group outcome standard deviation of 8. The resulting power is 0.87, as indicated in the last column of the table. The other columns present the results of additional LATE power simulations. For all of the simulations, the treatment effect and within-group variances for the compliers are also set to 5 and 8, respectively, while the compliance rate is set to 0.2. This implies that $\frac{100}{0.2^2} = 2500$ observations are required to reach a power of 0.87, and thus all simulations are set with a sample size of 2500. The results in the second and third columns of the table show that the power of the LATE estimator is indeed about 0.87 when the appropriate distributional assumptions are met with regard to the always-takers and never-takers (first specification). However, as the results also show, as soon as those assumptions are

violated, the scaled ATE power analysis provides extremely unreliable guidance on the power of the LATE estimator.

# Appendix C: Examples of Ordered-Means Assumption ($\bar{Y}_{NT} \leq \bar{Y}_C \leq \bar{Y}_{AT}$)

**Example 1: National JTPA Study (Bloom et al., 1997)**

Table C1: Example 1 Results by Treatment Assignment and Uptake

| | **D = 0** | **D = 1** |
|---|---|---|
| **Z = 0** | $N = 3663$<br>$\bar{Y} = 15062.99$ | $N = 54$<br>$\bar{Y} = 13515.26$ |
| **Z = 1** | $N = 2683$<br>$\bar{Y} = 13979.93$ | $N = 4804$<br>$\bar{Y} = 17439.8$ |

This example pertains to a well-known experimental study of the Job Training Partnership Act of 1982 (Bloom et al., 1997), where study subjects were randomly assigned such that they were either allowed to enroll in a job training program (assigned to the treatment) or excluded from the training for an 18-month period (assigned to control). The dataset used here is the tabulation of the JTPA study data by Abadie et al. (2002). The data correspond to adult participants in the JTPA experiment for whom there was a measurement of the outcome variable, which was earnings in the 30-month period after random assignment. Table C1 separates the sample into four groups based on treatment assignment ($Z$) and uptake ($D$). As shown in Table C1, there were almost no identified always-takers ($Z = 0, D = 1$) in the National JTPA Study. As explained in Lemma 6, the alternative power lower bound can be used when $Cov(E[D|X], E[Y|X])$ is non-negative, where $X$ denotes the principal strata, which is true given the ordered means assumption that $\bar{Y}_{NT} \leq \bar{Y}_C \leq \bar{Y}_{AT}$. However, given how few always-takers there are, they will have a negligible influence on $Cov(E[D|X], E[Y|X])$, and hence the ordered means assumption can be simplified to $\bar{Y}_{NT} \leq \bar{Y}_C$ as a practical matter.

There are 2683 identified never-takers ($Z = 1, D = 0$) in the study. Furthermore, this group of identified never-takers exhibits the lowest mean observed outcome among the three groups not including the identified always-takers. As a result, it is reasonable to conclude that the ordered means assumption was met in this study.

**Example 2: New Haven Vote Canvassing Experiment (Gerber and Green, 2000, 2005)**

Table C2: Example 2 Results by Treatment Assignment and Uptake

| | **D = 0** | **D = 1** |
|---|---|---|
| **Z = 0** | $N = 10733$<br>$\bar{Y} = 0.442$ | $N = 0$<br>$\bar{Y} = NA$ |
| **Z = 1** | $N = 1781$<br>$\bar{Y} = 0.435$ | $N = 869$<br>$\bar{Y} = 0.577$ |

This example pertains a voter mobilization field experiment conducted by Gerber and Green (2000) in 1998, using updated data from Gerber and Green (2005). The experiment involved three different types of voter mobilization treatments—mailings, phone calls, and home visits. The results presented here include those subjects assigned to receive no treatments (pure control) and those subjects assigned only to the home visits, as this was the treatment in which a major factor was noncompliance—i.e. not actually being contacted as a result of the subject not being home. Subjects for whom the outcome (having voted in the 1998 midterm election) was missing are omitted.

As Table C2 shows, there is strict one-sided noncompliance by design in this study, with no always-takers, thus simplifying the ordered means assumption to $\bar{Y}_{NT} \leq \bar{Y}_C$. The results are consistent with this assumption, as the identified never-takers ($Z = 1, D = 0$) have the lowest mean outcome value.

**Example 3: Fuzzy Regression Discontinuity Design on the Effect of Naturalization on Integration (Hainmueller et al., 2016)**

Table C3: Example 3 Results by Treatment Assignment and Uptake

| | $\mathbf{D=0}$ | $\mathbf{D=1}$ |
|---|---|---|
| $\mathbf{Z=0}$ | $N=69$<br>$\bar{Y}=-0.527$ | $N=90$<br>$\bar{Y}=0.048$ |
| $\mathbf{Z=1}$ | $N=0$<br>$\bar{Y}=NA$ | $N=244$<br>$\bar{Y}=0.055$ |

This example pertains to a study on the effects of naturalization on the political integration of immigrants who applied for Swiss citizenship prior to 2003 (Hainmueller et al., 2016),[10] during which time certain municipalities used secret ballot referendums to vote on which immigrants should be granted citizenship. As one of the identification and estimation strategies, the study includes a fuzzy regression discontinuity design comparing only those immigrants who barely received a majority "yes" vote and those who barely missed the majority. However, some of those immigrants who did not receive the majority vote initially went on to gain citizenship later. As a result, a fuzzy regression discontinuity design can be applied, where $Z$ indicates having received the "yes" vote (assignment to treatment) and $D$ indicates having gained citizenship (uptake of the treatment).

The fuzzy regression discontinuity design (RDD) combines the machinery of the sharp RDD and the LATE. The method of power analysis presented in this paper would be applicable to one particular estimator sometimes used with the fuzzy RDD, which is the standard LATE estimator applied to subjects within the specified RDD bandwidth. However, while the power method is technically not applicable to the actual fuzzy RDD regression estimators, it could be used as a convenient proxy.

As Table C3 shows, there is strict one-sided noncompliance by design in this study, with no never-takers, thus simplifying the ordered-means assumption to $\bar{Y}_C \leq \bar{Y}_{AT}$. The identified always-takers ($Z=0, D=1$) have almost the highest mean outcome value. However, the mean outcome of the compliers would be the weighted average of the mean outcome of the identified compliers in the ($Z=0, D=0$) group and the unidentified compliers in the ($Z=1, D=1$). While we cannot identify the specific units in the latter group that are compliers, rather

[10]The data are not publicly available due to privacy concerns, and the statistics in Table C3 were provided by the authors of the study. Individuals interested in accessing the data for replication purposes should contact the authors.

than always-takers, we have information on the identified always-takers ($Z = 0, D = 1$) and knowledge that the ($Z = 1, D = 1$) group is a mix of compliers and always-takers, which allows us to assess the ordered-means assumption.

Specifically, treating $Z$ as if randomized, we would expect the proportion of individuals who are always-takers to be equal across the ($Z = 0$) and ($Z = 1$) margins. Thus, we can estimate that approximately 138 of the individuals in the ($Z = 1, D = 1$) group are always-takers, whose mean outcome should be the same (in expectation) as the mean outcome of the identified always-takers in the ($Z = 0, D = 1$) group, which is 0.048. This implies that the mean outcome of the unidentified compliers in the ($Z = 1, D = 1$) group, of which we expect there should be 106, is approximately 0.064. As a result, the mean outcome of all compliers would undoubtedly be lower than the mean outcome of always-takers, given the particularly low value of the mean outcome for the identified compliers in the ($Z = 0, D = 0$) group. Specifically, using the information we have, we can estimate the mean outcome of all compliers to be $-0.17$. Thus, the results are consistent with the ordered-means assumption.

## Appendix D: Simulation Results Given $P(Z = 1) = 0.25$

Figure D1 compares simulated power curves to the analytic bounds given $P(Z = 1) = 0.25$, where power is plotted as a function of $\kappa$. As previously, the solid black lines denote the analytic upper and lower bounds, while the dashed black line denotes the alternative lower bound under the ordered-means assumption. The colored lines denote power curves that were simulated by specifying the full set of investigation and distribution parameters, and hence simulating all three principal strata. The specifications are the same as those pertaining to Figure 5 in the main text; the only difference is that $P(Z = 1) = 0.25$, rather than $P(Z = 1) = 0.5$.

Figure D1: Simulated Power vs. Analytic Bounds, $P(Z = 1) = 0.25$

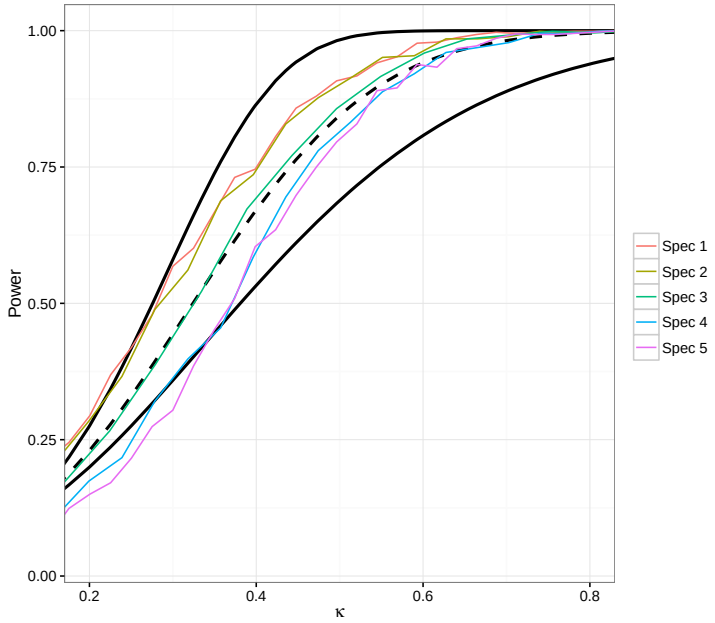


As Figure D1 illustrates, the analytic bounds derived for the general case where $P(Z = 1) = p_z$ perform as well as the bounds derived for the special case of $P(Z = 1) = 0.5$. Compared to Figure 5, the power curves (analytic and simulated) in Figure D1 are all slightly lower, which makes sense given that the decrease in $P(Z = 1)$ means a less balanced allocation of treatment uptake among the compliers. Similar to Figure 5, the alternative lower bound (the dashed black line) in Figure D1 bounds the appropriate simulated curves. For specifications 1 and 2, the ordered-means assumption is met at all values of $\kappa$, and hence the alternative lower bound applies. Accordingly, the curves for these specifications again lie above the alternative lower bound. In contrast, the ordered-means assumption is violated by specifications 4 and 5 at all

values of $\kappa$, and again these curves appropriately lie below the alternative lower bound.